

# High precision simulations of weak lensing effect on Cosmic Microwave Background polarization

Giulio Fabbian<sup>1\*</sup> and Radek Stompor<sup>1\*\*</sup>

AstroParticule et Cosmologie, Univ Paris Diderot, CNRS/IN2P3, CEA/Irfu, Obs de Paris, Sorbonne Paris Cité, France

March 27, 2013

## ABSTRACT

We study accuracy, robustness and self-consistency of pixel-domain simulations of the gravitational lensing effect on the primordial CMB anisotropies due to the large-scale structure of the Universe. In particular, we investigate dependence of the results precision on some crucial parameters of such techniques and propose a semi-analytic framework to determine their values so the required precision is a priori assured and the numerical workload simultaneously optimized. Our focus is on the  $B$ -mode signal but we discuss also other CMB observables, such as total intensity,  $T$ , and  $E$ -mode polarization, emphasizing differences and similarities between all these cases. Our semi-analytic considerations are backed up by extensive numerical results. Those are obtained using a code, nicknamed *lenS<sup>2</sup>HAT* – for Lensing using Scalable Spherical Harmonic Transforms (*S<sup>2</sup>HAT*) – which we have developed in the course of this work. The code implements a version of the pixel-domain approach of Lewis (2005) and permits performing the simulations at very high resolutions and data volumes, thanks to its efficient parallelization provided by the *S<sup>2</sup>HAT* library – a parallel library for a calculation of the spherical harmonic transforms. The code is made publicly available.

**Key words.** cosmic background radiation - large-scale structure of Universe - gravitational lensing: weak - methods: numerical

## 1. Introduction

The Cosmic Microwave Background (CMB) anisotropies in both temperature and polarization are one of the most studied signals in cosmology and one of the major available sources of constraints of the early universe physics. After having decoupled from matter and set free at the time of recombination, CMB photons propagated nearly unperturbed throughout the Universe. The large scale structure emerging in the universe in the post-recombination period has however left its imprint on them referred to as secondary anisotropies. In particular, the gravitational pull of the growing matter inhomogeneities has deviated the paths of primordial CMB photons modifying somewhat the pattern of the CMB anisotropies as observed today. This weak lensing effect on the CMB (see Lewis & Challinor (2006) for an extensive review) offers therefore a unique probe of matter distribution at intermediate redshift where the forming LSS was still in the nearly linear regime. As it depends on the cumulative matter distribution in the Universe it is expected to be particularly efficient in constraining the properties of all the parameters affecting the growth of Large Scale Structures (LSS), such as neutrino masses and dark energy physics (de Putter et al. 2009; Das & Linder 2012; Hall & Challinor 2012).

First observational evidence of the CMB lensing signal had been indirect and obtained through cross-correlation of the CMB maps with a high-redshift mass tracers (Smith et al. 2007; Hirata et al. 2008). More recently, its more direct measurements have become available, thanks to the latest generation of high precision and resolution ground based CMB temperature experiment, which have collected high quality data and made possible a direct reconstruction of the power spectra of those deviation using

CMB alone (Das et al. 2011; van Engelen et al. 2012). Even more recently, this has been further elaborated on by the Planck results based on the analysis of the first 15 months of the total intensity data as collected by the mission (Planck Collaboration 2013).

The forthcoming next generation of low noise CMB polarization experiment like EBEX (Oxley et al. 2004), POLARBEAR (Kermish et al. 2012), SPTpol (McMahon et al. 2009) and ACTpol (Niemack et al. 2010) and their future upgrades (Tomaru et al. 2012) will be able to target the CMB observable most affected by weak lensing. Primordial CMB gradient-like polarization ( $E$ -modes) is converted into curl-like polarization ( $B$ -modes) by gravitational lensing (Zaldarriaga & Seljak 1998) and thought to completely dominate the primordial signal at least at small angular scales. The lensing generated  $B$  modes are interesting due to their sensitivity to the large scale structure distribution, but also because they are the main contaminant of any primordial  $B$  modes signal, which is expected in many models of the very early universe, and which is one of the major goals of the current and future CMB observations. Since sensitivities of the CMB polarization arrays are rapidly improving, the experiments aiming at setting constraints on values of the tensor to scalar ratio parameter  $r \lesssim 10^{-2}$  are expected to be ultimately limited by lensing signal (e.g., Errard & Stompor 2012). The latter acts as an extra noise source with white spectrum shape on large scales and an amplitude of approximately  $5\mu K$ -arcmin, and which however could potentially be separated from the primordial signal with help of some accurate de-lensing procedure (Kesden et al. 2002; Seljak & Hirata 2004; Smith et al. 2012).

The high quality of forthcoming datasets requires the development, testing and validation through simulations of data analysis tools capable of fully exploiting the amount of information there present. An important part of this effort involves

\* e-mail: gfabbian@apc.univ-paris-diderot.fr

\*\* e-mail: radek@apc.univ-paris-diderot.fr

simulations of very accurate, high resolution maps of CMB total intensity and polarization, covering a large fraction of the sky and with lensing effects included. The relevant approaches have been studied in the past (e.g. Lewis 2005; Basak et al. 2009; Lavaux & Wandelt 2010) and resulted in devising and demonstrating an overall framework for such simulations, as well as two, publicly available, numerical codes (Lewis 2005; Basak et al. 2009). As computations involved in such a procedure are inherently very time consuming, the proposed implementations of those ideas involve unavoidably trade-offs between calculation precision and their feasibility, giving rise to a number of issues, practical and more fundamental, which need to be carefully resolved to ensure that such techniques produce high quality, reliable results. The main objective of this paper is to provide comprehensive answers to some of those problems, with special emphasis on those arising in the context of high precision and reliability simulations of the B-mode component of the CMB polarization signal.

## 2. Simulating weak lensing of CMB

### 2.1. Algebraic background

The CMB radiation is completely described by its brightness temperature anisotropies and polarization fields on the sky, respectively  $T(\vartheta, \varphi)$  and  $P(\vartheta, \varphi)$ . Since both fields are (nearly) Gaussian they are characterized by their power spectra after their Fourier expansion on a proper basis. Temperature is a scalar field and can be conveniently expanded in terms of scalar spherical harmonics

$$T(\vartheta, \varphi) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l T_{lm} Y_{lm}(\vartheta, \varphi) \quad (1)$$

while polarization depends on the Stokes parameters  $Q$  and  $U$ , which are base dependent vectors, and thus behaves like a spin-2 field on the sphere if rotated (Zaldarriaga & Seljak 1997; Kamionkowski et al. 1997). For this reason, polarization field must be expanded in terms of spin-2 spherical harmonics  ${}_{\pm 2}Y_{lm}(\vartheta, \varphi)$

$$P(\vartheta, \varphi) = (Q + iU)(\vartheta, \varphi) \quad (2)$$

$$= \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l -({}_2E_{lm} + i{}_2B_{lm}) {}_2Y_{lm}(\vartheta, \varphi)$$

where  ${}_2E_{lm}$  and  ${}_2B_{lm}$  are the gradient and curl harmonic components of a spin-2 field, whose general definitions for an arbitrary spin- $s$  field are

$${}_{|s|}E_{lm} \equiv -\frac{1}{2} \left( {}_{|s|}a_{lm} + (-1)^s {}_{-|s|}a_{lm} \right) \quad (3)$$

$$i{}_{|s|}B_{lm} \equiv -\frac{1}{2} \left( {}_{|s|}a_{lm} - (-1)^s {}_{-|s|}a_{lm} \right)$$

Weak gravitational lensing shifts the light rays coming from an original direction  $\hat{n}$  on the Last Scattering Surface to the observed direction  $\hat{n}'$ , inducing a mapping between the two directions through the so called displacement field  $\mathbf{d}$ , i.e. for a CMB observable  $X \in \{T, Q, U\}$

$$\tilde{X}(\hat{n}) = X(\hat{n}') = X(\mathbf{n} + \mathbf{d}). \quad (4)$$

Hereafter, we use a tilde to denote a lensed quantity, we will also use a tilde over a multipole number of a lensed quantity,

i.e.,  $\tilde{\ell}$ , to distinguish it from a multipole number of its unlensed counterpart.

The displacement field is a vector field on the sphere and can be decomposed into a gradient-free and a curl-free component. In most of the cases we can neglect the gradient-free component and consider the displacement field  $\mathbf{d}$  as the gradient of the so-called lensing potential  $\Phi(\vartheta, \varphi)$ , the projection of the 3-D gravitational potential  $\Psi$  on the 2-D unit sphere. Such quantity can be computed with Boltzmann codes (e.g. CAMB<sup>1</sup> or CLASS<sup>2</sup>), from galaxy surveys or N-Body simulations (Carbone et al. 2008; Das & Bode 2008)

$$\Phi(\mathbf{n}) \equiv -2 \int_0^{\eta_*} d_A \eta \frac{d_A(\eta_* - \eta)}{d_A(\eta) d_A(\eta_*)} \Psi(\eta, \mathbf{n}) \quad (5)$$

Here  $\eta_*$  is the comoving distance to the last scattering surface,  $\eta$  is the co-moving distance,  $d_A$  is the co-moving angular diameter distance. The lensing potential is expected to be correlated on large scale with temperature anisotropies and E modes of polarization through Integrated Sachs-Wolfe effect; this correlation affects mainly the large angular scales and is of the order of 1% at  $\ell \approx 100$  and will thus be neglected in the following analysis. Since the lensing potential is a scalar function and can be expanded into canonical spherical harmonics, the computation of its gradient (a spin-1 curl-free field) can be easily done in the harmonic domain with a spin-1 spherical harmonic transform (SHT).

$${}_1E_{lm} = \sqrt{l(l+1)} \Phi_{lm} \quad {}_1B_{lm} = 0 \quad (6)$$

### 2.2. Pixel-domain simulations

#### 2.2.1. Basics

Because typical deviations of CMB photons are of the order of few arcminutes (although coherent over the degree scale), we can work in the Born approximation, i.e. considering such deviation as constant between  $\hat{n}$  and  $\hat{n}'$ , and evaluate the displaced field along the unperturbed direction.

In practice this means that in order to compute the lensed CMB at a point it is sufficient to compute the unlensed CMB at another position on the sky. This observation provides the basis for the pixel-based approaches to simulating lensing effects of the CMB maps. For every direction on the sky corresponding to the pixel centers these methods first identify the displaced direction, and then compute the corresponding sky signal value, which is then used to replace the original value at the pixel center. The implementations of such approach typically follow these few main steps (Lewis 2005; Basak et al. 2009; Lavaux & Wandelt 2010):

1. Generation of a random realization of the harmonic coefficient of unlensed CMB map and synthesis of the field.
2. Generation of a random realization of harmonic coefficient of lensing potential and then of spin-1 displacement field in the harmonic domain. Synthesis of displacement field.
3. Sampling of displacement field at pixel centers and, for each of them, computation of the coordinates of displaced direction on the sky using the spherical triangle identities on the sphere.

Defining  $\alpha$  as an angle between the displacement vector and the  $\mathbf{e}_\vartheta$  versor, such that  $\mathbf{d} = d \cos \alpha \mathbf{e}_\vartheta + d \sin \alpha \mathbf{e}_\varphi$ , the value

<sup>1</sup> <http://camb.info>

<sup>2</sup> <http://lesgourg.web.cern.ch/lesgourg/class.php>

of a lensed field, i.e.,  $T$ ,  $Q$  and  $U$ , in a direction  $(\vartheta, \varphi)$  is given by the unlensed field at  $(\vartheta', \varphi + \Delta\varphi)$  where,

$$\cos \vartheta' = \cos d \cos \vartheta - \sin d \sin \vartheta \cos \alpha \quad (7)$$

$$\sin \Delta\varphi = \frac{\sin \alpha \sin d}{\sin \vartheta'} \quad (8)$$

4. Computing temperature and polarization fields at displaced positions.
5. Reassigning the temperature and polarization from the displaced to new positions to create the simulated lensed map sampled on the original grid. In the case of the polarization, we need also to multiply the lensed field by an extra factor taking into account the different orientation of the basis vector at the two points. Calling  $\gamma$  the difference between the angles between  $\mathbf{e}_\vartheta$  and the geodesic connecting the two points and defining

$$A = \tan \alpha' = \frac{d_\varphi}{d \sin d \cot \vartheta + d_\vartheta \cos d} \quad (9)$$

$$e^{2i\gamma} = \frac{2(d_\vartheta + d_\varphi A)^2}{d^2(1 + A^2)} - 1 + \frac{2i(d_\vartheta + d_\varphi A)(d_\varphi - d_\vartheta A)}{d^2(1 + A^2)} \quad (10)$$

the lensed polarization field becomes

$$\tilde{P}(\vartheta, \varphi) = e^{2i\gamma} P(\vartheta', \varphi') \quad (11)$$

6. Smoothing and, potentially, re-pixelizing the lensed map to match a particular experimental resolution if needed.

### 2.2.2. Challenges and paper goals

There are two main, closely intertwined challenges involved in an implementation of the approach detailed in the previous Section. The first one is related to the bandwidths of fields used in, or produced as a result of, the calculation, and in particular to the need of imposing those on the fields, which are either naturally not band-limited or are band-limited but with the band-widths too high to make them acceptable from the computational efficiency point of view. The other challenge arises on step 4 of the algorithm and is due to the fact that the displaced directions do not correspond in general to pixel centers of any iso-latitudinal grid on the sphere, and thus the lensed values of the CMB signal can not be computed with the aid of a fast SHT algorithm and a more elaborated, and computationally costly approach is needed.

We emphasize that both these problems should be looked at from the perspective of the efficiency of the numerical calculations as well as accuracy of the produced results. We discuss them in some detail below.

**Signal bandwidths.** As the lensing procedure needs to be applied prior to any instrumental response function convolution, the relevant sky signals on all, but the last, steps above require using a resolution sufficient to support the signal all the way to its intrinsic bandwidth,  $\ell_{intr}^X$ , where  $X$  is either  $T$  for the total intensity, or  $P$  for the polarization, or  $\Phi$  – for the gravitational potential. However, as mathematically the lensing effects can be seen as a convolution in the harmonic domain (Hu 2000; Okamoto & Hu 2003; Hu & Okamoto 2002) of the CMB signal – either the total intensity,  $T$ , or the polarization,  $P$ , – and of the potential,  $\Phi$ , the bandwidth of the resulting lensed field will be larger than that of any unlensed fields and given roughly by  $\ell_X^{intr} + \ell_\Phi^{intr}$ .


Consequently, the lensed map produced on step 5 should have its resolution appropriately increased to eliminate potential power aliasing effects. The resolution of the unlensed maps produced on steps 1-5 should then coincide with that of the lensed signal but with the number of harmonic modes as set respectively by  $\ell_X^{intr}$  and  $\ell_\Phi^{intr}$ .

One of the problems arising in this context is related to the fact that the sky signals,  $T$ ,  $P$  and  $\Phi$ , considered here are not truly band-limited even if their power at the small scales decays quite abruptly as a result of Silk damping. Picking an appropriate value for the bandwidth is therefore a matter of a compromise between the precision of the final products and the calculation cost, with both these quantities being quite sensitive to the chosen value, and which will depend in general on a specific application. We emphasize that the presence of the high- $\ell$  power decay plays a dual role in our considerations here. On the one hand, it ensures that the lensing effect at sufficiently large scale can be computed with an arbitrary precision, by simply choosing the bandwidth values sufficiently large. On the other hand it does introduce an extra complexity in defining a set of sufficient conditions, which ensure required precision, as these will be typically different in the regime of the high signal power and that of the damping tail. In either case though, it is clear that whatever are the selected bandwidths, the amplitudes of the harmonic modes close to the highest value of  $\tilde{\ell}^X$  supported by the employed pixelization, i.e.,  $\tilde{\ell}^X \sim \ell_{intr}^X + \ell_{intr}^\Phi$ , will generally be unavoidably misestimated, while satisfactory precision could only be achieved for sufficiently low harmonic modes,  $\tilde{\ell}^X < \tilde{\ell}_{ok}^X < \ell_{intr}^X$ . From the practitioner's perspective the major issue is therefore how given some precision criterion,  $\varepsilon$ , which we wish to be fulfilled up to some value of  $\tilde{\ell}^X = \tilde{\ell}_{ok}^X$ , to determine the required bandwidths,  $\ell_{intr}^X = \ell_{intr}^X(\tilde{\ell}_{ok}^Y, \varepsilon)$  where  $X$  and  $Y$  can be the same, e.g., in the case of the  $T$  or  $E$  signal lensing, or different, e.g., for the potential field or  $B$ -modes.

One upshot of these considerations is that if these are maps which are of interest as the final product of the calculation, then the biased high- $\ell$  modes, should be either filtered out or suppressed, before the map is synthesized from its harmonic coefficients. To ensure that this does not affect adversely the resolution of the final map, the bias should affect only angular scales much smaller than the expected final resolution of the map as produced on step 6 of the algorithm. If the latter is defined by the experimental beam resolution, one therefore needs to ensure that no bias is present for  $\tilde{\ell}^X \lesssim \ell_{beam} \sim \sigma_{beam}^{-1}$ , where  $\sigma_{beam}$  is an experimental beam width.

**Interpolations.** Interpolation is the most popular workaround of the need to calculating directly the values of the unlensed fields for every displaced directions, which typically will not correspond to grid points of any iso-latitudinal pixelization. Three interpolation schemes have been considered to date in the context of the polarized signals. Lewis (2005) proposes a generic modified bicubic interpolation and demonstrates that it seems to work satisfactory in a number of cases. This approach together with the direct summation are both implemented in the publicly available code LensPix<sup>3</sup>. Two other methods have been proposed more recently. Basak et al. (2009) implements the general interpolation scheme, which recast a band-limited function on the sphere as a band-limited function on the 2-D torus where a non-equispaced Fast Fourier (NFFT) transform algorithm is used to compute the field at the displaced positions. This method would

<sup>3</sup> <http://cosmologist.info/lenspix>



**Fig. 1.** Examples of the CMB B-mode lensing calculation and involved numerical effects. All the panels show the recovered  $B$ -mode power spectrum, color lines, overplotted over the theoretical  $B$ -mode spectrum computed with CAMB. The bandwidth of the  $E$ -mode and the potential  $\Phi$  is the same in all the panels and set to 2500, while the resolution of the maps used for simulating the lensed signal increases progressively from left to right. ECP pixelization has been used in all the cases. The recovered  $B$ -spectrum overestimates the theoretical curve in the left panel due to the power aliasing effect, while it underestimates it in the result recovered for much higher resolution as shown on the right. The nearly perfect recovery shown in the middle panel is merely accidental and results from the extra contribution due to aliasing (left panel) being compensated by the insufficient signal bandwidth (right panel). The spectrum in the right panel is aliasing free as it does not change anymore with the increasing resolution.

have been arbitrarily precise if the sky signals were strictly band-limited. Nevertheless the choice of NFFT can become a bottleneck for this algorithm since its numerical workload scales with the number of pixels squared, and its memory requirements are huge. We note that, as is, the NFFT software can be run only on shared memory architectures, making more difficult to resolve both of these issues. Consequently, the relevance of the approach in particular in the context of simulations of upcoming and future high resolution experiments needs to be further investigated. In any case, the issue of the bandwidth values is becoming of crucial importance for the performance and applicability of the method.

Lavaux & Wandelt (2010) propose a fast pixel-based method using the spectral characteristics of the field to be lensed to compute the weighting coefficient for the interpolation of such field, without using any spherical harmonic algorithm. Its accuracy is set by the number of neighboring pixel used to interpolate the field at a given point.

In addition, Hirata et al. (2004) use in their work a polynomial interpolation scheme of arbitrary order and precision, which has been shown to produce successfully temperature maps (Hirata et al. 2004; Das & Bode 2008) but not tested for the polarized case.

Any interpolation in this context is not without its dangers as interpolations tend to smooth the underlying signals. For genuinely band-limited function that could be in principle avoided as in e.g., Basak et al. (2009). However, for the actual CMB signals the band-width is only approximate and is a function of the required precision and specific application and the sampling density and interpolation scheme need to be chosen very carefully to render reliable results. Again the choice of the appropriate band-width values is therefore central for a successful resolution of this problem.

**Numerical workload** Numerical cost of the direct calculation per direction is given by  $O(\ell_{max}^2)$  and it corresponds to the cost of calculating an entire set of all  $\ell$  and  $m$  modes of associated, scalar or spin-weighted, Legendre functions. For  $N_{pix}$  direction the overall cost is on order of  $O(N_{pix} \ell_{max}^2) = O(N_{pix}^2)$  and is therefore prohibitive for any values of  $N_{pix}$  and  $\ell_{max}$  of interest. Here, we have assumed a relation,  $\ell_{max} \propto N_{pix}^{1/2}$ , typically

fulfilled for the full sky pixelization with a proportionality coefficient on order of a few, e.g., for the HEALPix<sup>4</sup> pixelization (Górski et al. 2005) we have  $\ell_{max} = 2\sqrt{3N_{pix}}$ , while for ECP,  $\ell_{max} = 2\sqrt{N_{pix}}$ . The interpolations can cut on this load trimming it to the one needed to compute a representation of the signals on a iso-latitudinal grid, what has the complexity of  $O(N_{pix}^{1/2} \ell_{max}^2) = O(N_{pix}^{3/2})$  plus the interpolation with the complexity  $O(N_{pix})$ , or  $O(N_{pix} \ln N_{pix})$  in the case of NFFT, in both cases with a potentially large pre-factor. Nevertheless, this is clearly a more favorable scaling than the one of the direct method and one, which as has been shown in the past, makes such calculations feasible in practice. We note however that for the sake of the precision of the interpolation one may need to overpixelized the sky, meaning using a higher value of  $N_{pix}$  than what would normally be needed to support the harmonic modes all the way up to  $\ell_{max}$ . Hereafter, we will denote the overpixelization factor in one, i.e., either  $\theta$  or  $\phi$  direction, as  $\kappa$ . Consequently the number of pixels used is given by  $\kappa^2 N_{pix}$ , where  $N_{pix}$  is the standard full sky number of pixels as determined by the selected value of  $\ell_{max}$ .

**Paper goals and methodology.** This paper has two main goals. One is to study internal consistency and convergence of the pixel-domain simulations in the context of the currently viable cosmologies. The other is to study the dependence of the precision of such simulations on some of its most important parameters.

We note that, in the previous works, analyses of this sort have been usually restricted to comparisons of power spectra of the lensed maps derived by a lensing simulation code and the theoretical predictions computed via an integration of the Boltzmann equation, as implemented in the publicly available codes, CAMB and CLASS. In these works, the effort has been made to find a set of the codes parameters for which the resulting spectrum is consistent with the theoretical expectations. Such comparisons are without doubt an important part of the code and method validation. However, they are limited to the cases of the gravitational potentials,  $\Phi$ , derived in a linear theory, and not applicable in some other cases where the potential

<sup>4</sup> <http://healpix.sf.net/>



is obtained by some other means as for instance, N-body simulations. In addition, they may be on occasion misleading as the numerical effects can easily conspire to deliver a spectrum tantalizingly close to the desired one, without any reassurance that the map of the lensed sky characterized by it, has correct other statistical properties, such as e.g., higher order statistics. We note that this is particularly likely and consequential for the  $B$ -mode spectra given their low amplitudes and the lack of characteristic, fine-scale features. An example of such a conspiracy is shown in Fig. 1, where the power deficit at the high- $\ell$  end caused by the oversmoothing due to the interpolation compensate nearly perfectly the extra power aliased into the  $\ell$ -range of interest as a consequence of a too crude resolution of the final map.

In this work we therefore propose to study the robustness of the simulated results by demonstrating their convergence and internal stability with respect to sky sampling and band-limit changes, as expressed by two parameters, introduced earlier: the upper value of the signal band,  $\ell_{max}$ , and the overpixelization factor,  $\kappa$ . Only once the convergence is reached we compare the results to those computed by other means, if such are available. We note that the convergence tests do not have to, and should not in general, be restricted to the power spectra comparison only and could instead involve other metrics more directly relevant to the simulated maps themselves. In all such tests it is typically required to consider maps with extreme resolutions, what has been traditionally prohibitive for numerical reasons. In this work, we overcome this problem with help of a high performance lensing code, `lenS2HAT`, which we have developed for this purpose.

Our second goal, i.e., to study the dependence of the calculation precision on the two crucial parameters,  $\ell_{max}$  and  $\kappa$ , is complementary and is aimed at providing meaningful and practically useful guidelines of how to select the values of these parameters prior to performing any numerical tests given some predefined precision targets. In this context, we present here an in-depth semi-analytical analysis of the impact of these parameters on the lensed signal recovery. Though ultimately they may need to be confirmed numerically case-by-case, e.g., using the convergence tests as discussed earlier, they could be of significant help in providing a reasonably starting point for such tests.

We also discuss a simple, high performance parallel implementation of the pixel-domain lensing code, `lenS2HAT`, capable of reaching extremely high sample density on the sphere, thanks to its efficient parallelization and numerical implementation, and which has been instrumental in accomplishing all the other goals of this work. The code is made publicly available.

### 3. Exploring the bandlimits

#### 3.1. CMB lensing in harmonic domain

This Section addresses the second of the goals of this paper, as stated above, and it describes a semi-analytic study of the impact of the assumed bandwidth values on the precision of the lensed signal. Our discussion is based on the model of Hu (2000) and focuses on the case of the lensed  $B$ -mode signal as obtained as a result of the lensing acting upon the  $E$ -mode one, which is the main target of this paper. Similar considerations can be however done for other CMB observable spectra and we will present some relevant results calculated also for those cases. (See Sect. 3.2 for some more details.) Using the results of Hu (2000) we represent the lensed  $B$ -mode signal as,

$$\tilde{C}_{\tilde{\ell}^B}^{BB} = \frac{1}{2} \sum_{\ell^\Phi \ell^E} \frac{|2F_{\tilde{\ell}^B \ell^\Phi \ell^E}|^2}{2\tilde{\ell}^B + 1} C_{\ell^\Phi}^{\Phi\Phi} C_{\ell^E}^{EE} (1 - (-1)^L) \quad (12)$$

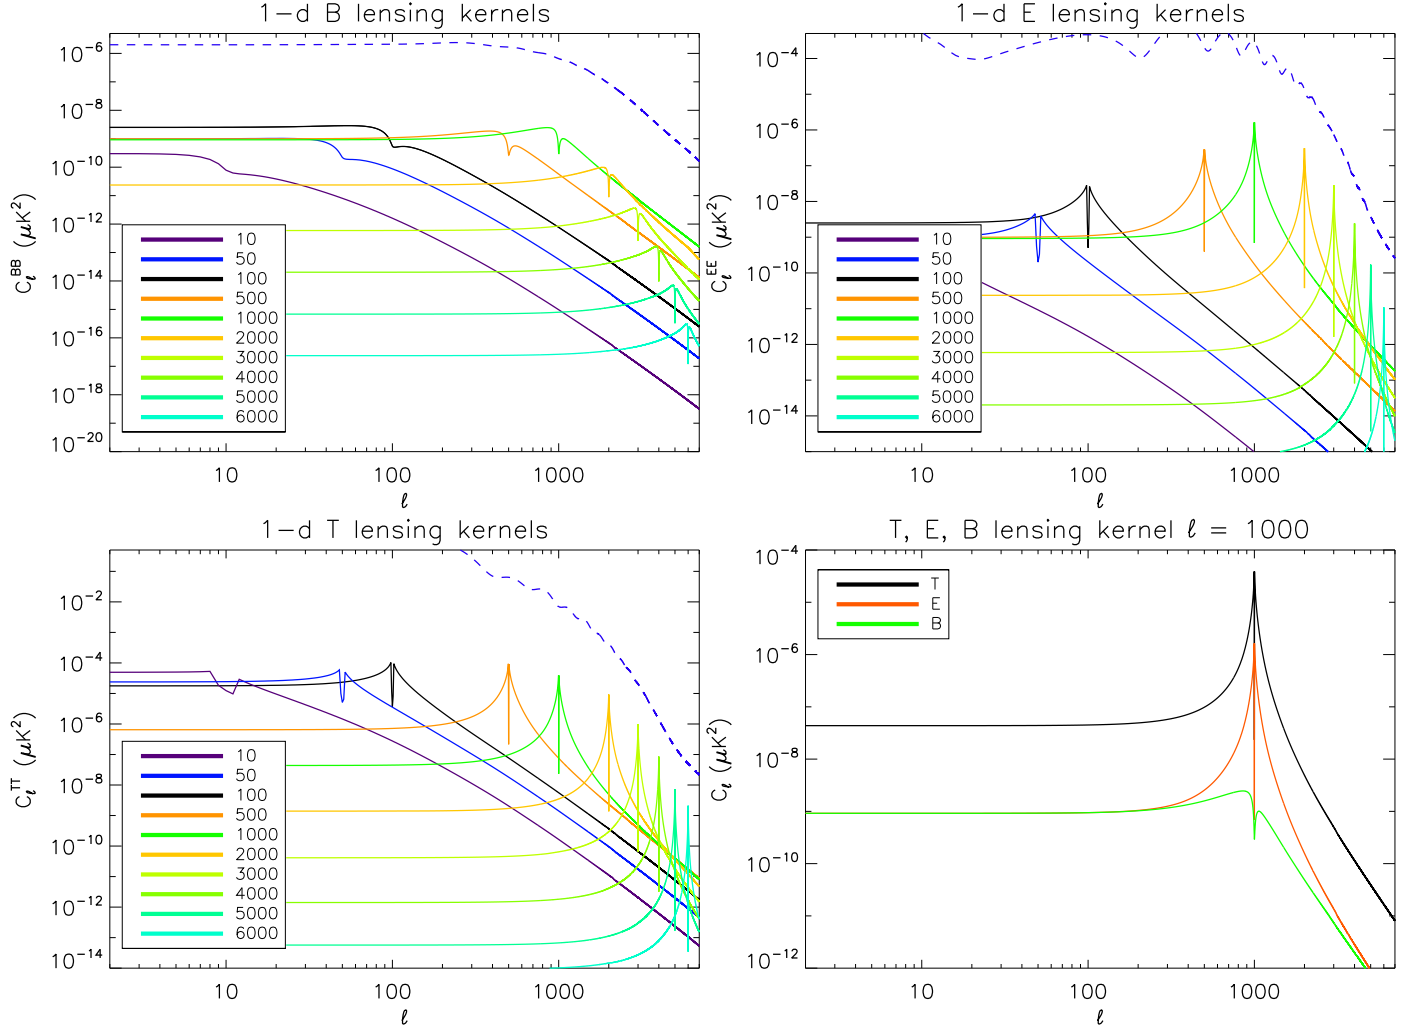
where  $2F_{\tilde{\ell}^B \ell^\Phi \ell^E}$  is a spin-2 coupling kernel (see Hu (2000) for a full expression),  $L \equiv \tilde{\ell}^B + \ell^\Phi + \ell^E$  and  $C_{\ell^E}^{EE}$  and  $C_{\ell^\Phi}^{\Phi\Phi}$  denote the unlensed power spectra of the  $E$  mode polarization and of the gravitational potential, respectively. This formula can be obtained by a second order series expansion around undispersed direction and it is expected to be accurate to within 1% for multipoles  $\tilde{\ell}^B \lesssim 2000$  and then for  $\tilde{\ell}^B \gg 2000$ , where the CMB amplitude is small and can be modeled by its gradient only, while in the intermediate scales its precision degrades to nearly 5%. The reliability of this analytical model will be discussed later in Sec. 4.3.1. We can now introduce 1-dimensional kernels,  $\mathcal{H}_{\ell^E}(\tilde{\ell}^B)$  defined as,

$$\mathcal{H}_{\ell^E}(\tilde{\ell}^B) \equiv \frac{1}{2} C_{\ell^E}^{EE} \sum_{\ell^\Phi} \frac{|2F_{\tilde{\ell}^B \ell^\Phi \ell^E}|^2}{2\tilde{\ell}^B + 1} C_{\ell^\Phi}^{\Phi\Phi} (1 - (-1)^L). \quad (13)$$

and which summed over  $\ell^E$  for a fixed  $\tilde{\ell}^B$  give the lensed  $B$ -mode power contained in the mode  $\tilde{\ell}^B$ , Eq. 12, while for a fixed  $\ell^E$  they define power spectrum of the lensed  $B$  mode signal, generated via lensing from the  $E$  polarization signal containing non-zero power in a single mode  $\ell^E$ , and with its amplitude as given by  $C_{\ell^E}^{EE}$ . The kernels are displayed in Fig. 2 together with their analogs for the total intensity and  $E$ -polarization signals. We find that the kernels computed for different values of  $\tilde{\ell}^B$  are similar just shifted with respect to each other accordingly. The change in the amplitude reflects simply the change in the assumed power of the  $E$  signal, which in turn follows that of the actual  $E$  power spectrum. The kernels are flat for values  $\tilde{\ell}^B \ll \ell^E$  and decay as a power law for  $\tilde{\ell}^B \gg \ell^E$ , displaying a sharp dip at  $\tilde{\ell}^B = \ell^E$ . Similar observations can be made for the  $T$  and  $E$  kernels, with an exception that, unlike their  $E$  and  $T$  counterparts, the  $B$  kernels are not peaked around that value. Such behavior is related to the fact that the lensed  $B$  mode signal we discuss here, and as described by Eq. 12, is due to the  $E$ -polarization, while the main effect of the lensing on  $T$  and  $E$  is imprinted on these signals directly. A direct consequence of this fact is that for any lensed  $B$  spectrum mode a contribution from local unlensed multipoles will be less dominant, as it is the case for the  $T$  and  $E$  signals, and non-local contributions will be relatively more important and therefore required to be accounted for high precision calculations.

Indeed, due to the flat plateau of the kernels at the low- $\ell$  end, in principle all high- $\ell$  unlensed modes contribute to the lensed power at the low- $\ell$  end. The magnitude of their contribution is modulated by the shape of the unlensed  $E$  spectrum and therefore eventually becomes negligible only due to the presence of the Silk damping, i.e., lack of the power at small angular scales in the unlensed fields. Nevertheless, we can expect that nearly all the modes of the unlensed  $E$  spectrum up to the damping scale have to be included in the calculation of the lensed  $B$  spectrum to ensure high precision recovery of the lensed  $B$  spectrum with  $\tilde{\ell}^B \lesssim 1000$ . Given some specific target precision, we could and should fine-tune the required  $E$ -spectrum bandwidth and whatever is the value selected here the bandwidth for the potential field will have to be at least the same.

For the high  $\ell$  modes of the lensed  $B$  spectrum,  $\tilde{\ell}^B \gg 1000$  the non-locality of the power transfer due to lensing is even more striking, as due to the low amplitudes of the  $E$  spectrum the local contributions are additionally suppressed, and the long power-law tails of the contributions from large and intermediate angular scales,  $\ell^E \lesssim 1000$  are evidently dominant. Less evident is the fact that also the  $E$ -power from even smaller angular scales,  $\ell^E \gtrsim \tilde{\ell}^B$ , may be relevant. The contributions from each of those modes may appear small, Fig. 2, but are potentially non-




**Fig. 2.** 1-dimensional lensing kernels. The lensed power for T, E, B spectra is computed assuming a delta-like spectra with power in a single mode  $\ell' = 10, 50, 100, 500, 1000, 2000, 3000, 4000, 5000, 6000$  in the unlensed CMB spectra. The blue dashed line represents the reference lensed spectra as computed by CAMB. The sum of all single modes contribution for  $\ell' \in [0, \infty]$  would reproduce the lensed spectra. For T and E cases, the subdominant contribution of convolution part only is shown for visualization purpose and offset term are ignored (see Sect. 3.2 and Eq. 20). The comparison of 1-dimensional kernels shapes for T, E, B for  $\ell' = 1000$  is shown in the bottom right panel: the peculiar shape of each kernel type drives the locality and amplitude of contribution to the lensed spectra.

negligible due to a large number of those modes. A high precision recovery of the high- $\ell$  tail of the lensed  $B$  mode spectrum will therefore need a careful assessment of the importance of all these contributions, nevertheless a generic expectation would be that the bandwidth of the unlensed  $E$  spectrum will have to be higher than the maximal value of the lensed  $B$  signal multipole for which high precision is required and potentially higher than Silk damping scale. As these very high multipoles of the lensed  $B$  spectrum are expected to have a significant contribution coming from relatively low multipoles of the unlensed  $E$  signal, i.e., for which  $\ell^E \ll \tilde{\ell}^B$  given the triangular relations, Eq. 16, and the definition of the kernels, Eq. 13, we can conclude that the bandwidth of the potential field used in the simulations will have to be at least as large as  $\tilde{\ell}^B$ .

There are two main conclusions to be drawn here. First, it is clear that a high fidelity simulation of the  $B$ -polarization power spectrum even in a restricted range of angular scales will require large bandwidths, potentially all the way up to the scale of Silk damping, for both the unlensed  $E$ -mode polarization sig-

nal and the gravitational potential. These bandwidths values are not however expected to depend very strongly on the maximal  $B$ -mode multipole to be recovered precisely at least as long as it is in the range  $\tilde{\ell}^B \lesssim 2000$ . Second, as the expected bandwidths are large, it is important to optimize them in order to ensure efficiency of the numerical codes, without affecting precision of the results.

For the lensed  $T$  and  $E$  spectra, thanks to the peaked character of the respective kernels, the lensed modes are typically dominated by the local contribution coming from immediate vicinity of the mode. This permits in general setting the bandwidth for the potential shorter than the mode of the lensed spectrum to be computed. By contrast the unlensed  $T$  and  $E$  spectrum have to be known at least up to the multiple of interest of the lensed spectrum,  $\tilde{\ell}^X$ , ( $X = T$  or  $E$ ), augmented by the assumed bandwidth of the potential. These observations reflect the usual rule of thumb, (e.g., Lewis 2005), indicating that lower bandwidth values can be used in these two cases for the same required accuracy.



**Fig. 3.** Lensing kernels  $\mathcal{K}_{\tilde{\ell}^X}(\ell^Y, \ell^\Phi)$  for  $X = Y = T$ , left column,  $X = Y = E$ , middle, and  $X = B, Y = E$ , right, and for two different values of the multiple number of the lensed signal,  $\tilde{\ell}^X = 500, 1000$ , top to bottom. The color scale shows the logarithm of the kernel elements and ranges from dark blue  $\sim 10^{-15}$  to  $\sim 1$ , dark red. The solid line contours show the best achievable precision of the estimated lensed spectrum, which can be obtained if the bandwidths of the  $E$  and/or  $\Phi$  unlensed spectra are truncated to  $\ell^E$  and  $\ell^\Phi$ . The contours range from 25% down to 0.01% from left to right. The precision is computed with respect to the lensed multipoles calculated with  $\ell_{max}^E = \ell_{max}^\Phi = 8000$ .

### 3.2. Accuracy

In this Section we aim at turning the consideration presented above into more quantitative prescriptions concerning the bandwidths of the input fields used in the simulations. For this reason we introduce 2-dimensional kernels,  $\mathcal{K}_{\tilde{\ell}^B}(\ell^E, \ell^\Phi)$ , defined as,

$$\mathcal{K}_{\tilde{\ell}^B}(\ell^E, \ell^\Phi) \equiv \frac{1}{2} \frac{|{}_2F_{\tilde{\ell}^B \ell^\Phi \ell^E}|^2}{2\tilde{\ell}^B + 1} C_{\ell^\Phi}^{\Phi\Phi} C_{\ell^E}^{EE} (1 - (-1)^L). \quad (14)$$

These define for a given value of  $\tilde{\ell}^B$  a contribution of the  $E$  power at  $\ell = \ell^E$  and the  $\Phi$  power at  $\ell = \ell^\Phi$  to the amplitude of the lensed  $B$  mode spectrum at that  $\ell = \tilde{\ell}^B$ , which can be then computed by summing over  $\ell^E$  and  $\ell^\Phi$ , i.e.,

$$\tilde{C}_{\tilde{\ell}^B}^{BB} = \sum_{\ell^\Phi, \ell^E} \mathcal{K}_{\tilde{\ell}^B}(\ell^E, \ell^\Phi). \quad (15)$$

The sum in this equation involves in principle an infinite number of terms and therefore would have to be truncated in any numerical work, either explicitly, e.g., by setting finite limits in the formula above, or implicitly, e.g., by selecting the bandwidths, pixel sizes, etc., in the pixel-domain codes. We will therefore use these kernels to study the precision issues involved in this kind of calculations. As the expressions for the kernels are approximate, so will be our conclusions. However, as our goal is to


provide guidelines how to select the right values for the simulations codes, this should not pose any problems. We will get back to this point later in this Section.

We show a sample of the kernels,  $\mathcal{K}_{\tilde{\ell}^B}(\ell^E, \ell^\Phi)$  in Fig. 3. Those are computed for selected values of  $\tilde{\ell}^B$  for which the approximations involved in their computation are expected to be valid. We note that all elements of the kernel,  $\mathcal{K}_{\tilde{\ell}^B}(\ell^E, \ell^\Phi)$ , for which the quantity  $L$ , defined in the previous section, is even vanish, as do those for which the triangular relation,

$$|\ell^E - \ell^\Phi| \leq \tilde{\ell}^B \leq \ell^E + \ell^\Phi, \quad (16)$$

is not satisfied. This latter fact is a consequence of the presence of the Wigner 3-j symbols in the expressions for  ${}_2F_{\tilde{\ell}^B \ell^\Phi \ell^E}$ , (Hu 2000). Within these restrictions it is apparent from Fig. 3 that each multipole of lensed  $B$  mode spectra  $\tilde{\ell}^B$  receives contributions from a wide range of harmonic modes of both  $E$  and  $\Phi$  spectra, extending to values of  $\ell^E$  and  $\ell^\Phi$  significantly larger than  $\tilde{\ell}^B$  and roughly independent on the latter value at least for  $\tilde{\ell}^B \lesssim 2000$ . For its larger values a non-negligible fraction of the contribution starts coming from progressively higher multipoles of both  $E$  and  $\Phi$ . Clearly, these trends are consistent with what we have inferred earlier with help of the 1-dim kernels.

Also as observed earlier, we find the  $B$  mode kernels qualitatively different from those computed for the lensed total intensity and  $E$  mode polarization signals, Fig. 3 and are more localized in



**Fig. 4.** Lensing kernels  $\mathcal{K}_{\tilde{\ell}^X}(\ell^Y, \ell^\Phi)$  for  $X = Y = T$ , left column,  $X = Y = E$ , middle, and  $X = B, Y = E$ , right, and for two different values of the multiple number of the lensed signal,  $\tilde{\ell}^X = 2000, 4000$ , top to bottom. See Fig. 3 for additional details.

the harmonic space with the bulk of power coming mainly from scales, for which both  $\ell^{T,E}$  are relatively close to the considered lensed multipole,  $\tilde{\ell}^{T,E}$ .

We note that all the 2-dim kernels are positive<sup>5</sup> and therefore including more terms in the sum, Eq. 15, will always improve the result precision. From the efficiency point of view one may want to include in the sum preferably the terms corresponding the largest 2-dim kernel amplitudes as they provide the largest contribution to the final lensed result, before adding those with progressively smaller and smaller kernel amplitudes until the required precision is reached. Such an approach would in principle ensure that the best accuracy is achieved and done so with the smallest number of included terms. This may therefore look as a potentially attractive option from the perspective of the optimization of the calculations. However, in practice as the recurrence formulae are usually employed in the calculations, e.g., either those needed to compute spherical harmonics in the case of the pixel-domain codes or those needed to calculate the 3-j symbols as in a direct application of Eq. 15, and therefore all the terms up to a given bandwidth are at our disposal at any time and it therefore seems efficient and useful to capitalize on those by including all of them in the calculation. Consequently, we will be estimating what precision can be achieved by such calculations by including all the contributions up to some specific bandwidth values for the  $E$  and  $\Phi$  multipoles.

In the case of the  $B$ -spectra we therefore express hereafter the precision of the calculations as,

$$A_{\tilde{\ell}^B}^B(\ell^\Phi, \ell^E) = 1 - \frac{\sum_{\ell_*^\Phi=0}^{\ell^\Phi} \sum_{\ell_*^E=0}^{\ell^E} \mathcal{K}_{\tilde{\ell}^B}(\ell_*, \ell_*)}{\sum_{\ell_*^\Phi=0}^{\ell_{\max}^\Phi} \sum_{\ell_*^E=0}^{\ell_{\max}^E} \mathcal{K}_{\tilde{\ell}^B}(\ell_*, \ell_*)} \quad (17)$$

where the sums in the denominator should in principle extend over the infinite range of values of  $\ell$  but for practical reasons are truncated to  $\ell_{\max} = 8000$ , which for the range of lensed multipoles of interest in this work,  $\tilde{\ell}^X \lesssim 5000$ , should be sufficient.

This expression can be generalized for all the lensed CMB spectra but in this case our model has to take into account that main effect due to lensing is to reshuffle the power of the signal rather than convert it into some other component and therefore the total variance of the signal has to be conserved (e.g. Blanchard & Schneider 1987). In this case the lensed power spectra of  $X = T$  or  $E$  can be written as

$$\tilde{C}_{\tilde{\ell}^X}^X = (1 - (\tilde{\ell}^{X^2} + \tilde{\ell}^X - \alpha)R) C_{\tilde{\ell}^X}^X + \sum_{\ell^X, \ell^\Phi} \mathcal{K}_{\tilde{\ell}^X}(\ell^X, \ell^\Phi) \quad (18)$$

$$R = \sum_{\ell^\Phi}^{\ell_{\max}^\Phi} \frac{(\ell^\Phi + 1)(2\ell^\Phi + 1)}{4\pi} C_{\ell^\Phi}^\Phi \quad (19)$$

where  $\alpha$  is an integer which is different for each CMB spectra

- $\alpha = 2$  for  $X=E$
- $\alpha = 0$  for  $X=T$

<sup>5</sup> This is not true for the  $TE$  kernels, which we comment about later.

–  $\alpha = 1$  for  $X=TE$

We note that the factor  $R$  is a smooth function of the cutoff value of the sum over  $\ell^\Phi$ , which quickly becomes nearly constant for  $\ell_{max}^\Phi \gtrsim 1000$ , Fig. 5. Hereafter, we will therefore precompute it once assuming  $\ell_{max}^\Phi = \ell_{max} = 8000$  and use it in all subsequent calculations. The generalized expression for the accuracy function in Eq. 17 would then be


$$A_{\tilde{\ell}^X}^X(\ell^\Phi, \ell^X) = 1 - \frac{O_{\tilde{\ell}^X}^X + \sum_{\ell_*^\Phi=0}^{\ell^\Phi} \sum_{\ell_*^X=0}^{\ell^X} \mathcal{K}_{\tilde{\ell}^X}(\ell_*^X, \ell_*^\Phi)}{O_{\tilde{\ell}^X}^X + \sum_{\ell_*^\Phi=0}^{\ell_{max}^\Phi} \sum_{\ell_*^X=0}^{\ell_{max}^X} \mathcal{K}_{\tilde{\ell}^X}(\ell_*^X, \ell_*^\Phi)}, \quad (20)$$

where for shortness we have introduced,

$$O_{\tilde{\ell}^X}^X \equiv \left(1 - (\tilde{\ell}^{X^2} + \tilde{\ell}^X - \alpha)R\right) C_{\tilde{\ell}^X}^X.$$

We note that for cosmological models of the current interest, the factor,  $R$  is found typically to be on order of  $\mathcal{O}(10^{-6})$  and thus the term  $O_{\tilde{\ell}^X}^X$  is expected to be negative for most of the values of  $\tilde{\ell}^X$  in the range of interest here, Fig. 5.

In Fig. 3 black solid lines represent the expected estimation of



**Fig. 5.** Example of behavior of the term offset term  $O_{\tilde{\ell}^T}^T$  as a function of  $\ell_{max}^\Phi$  for  $\tilde{\ell}^T = 500$  (red), 2000 (blue), 4000 (green). Dashed line represent a negative value. The  $O_{\tilde{\ell}^E}^E$  and  $O_{\tilde{\ell}^B}^B$  have an analogous behavior but a different amplitude.

errors, as expressed by the accuracy function,  $A_{\tilde{\ell}^X}^X(\ell^\Phi, \ell^Y)$ , for a number of selected values ranging from 25% down to 0.01%. We note that for the shown range of  $\tilde{\ell}$  only the sub-percent values of the accuracy are likely to be somewhat biased due to the assumed cutoff in the denominator of Eqs. 17 or 20, which in all the cases can be considered at least indicative. On the other hand, the fact that our accuracy definition is based on an approximate formula is less of an issue as any potential (and small, Challinor & Lewis (2005)) deviations affect both the numerator and denominator of Eqs. 17 and 20 in the same way. It can be therefore shown that to the first order in the deviations amplitude, our accuracy criterion precision becomes progressively better and better, when the estimated level of the accuracy,  $A_{\tilde{\ell}^X}^X(\ell^\Phi, \ell^X)$ , tends to 0 and is degraded to the percent level when  $A_{\tilde{\ell}^X}^X(\ell^\Phi, \ell^X) \approx 90\%$ , i.e., when it is well outside of the region of any interest for high precision simulations

considered here (see appendix A).

The differences in the shape of the lensing kernels result in differences in the accuracy contours for different lensed signals and their multipoles. In particular, for lensed  $B$  modes, the contribution of large scale power of the CMB to the lensed signal is more significant. Nevertheless, the majority of the contours seem to share the same rectangular shape. Consequently, for any fixed value of one of the bandwidths, for instance,  $\ell^\Phi$  in the bottom left panel of Fig. 3, the achievable maximal accuracy can be already reached – at least for all the practical purposes – for the other bandwidth, e.g.,  $\ell^E$ , set to a value,  $\ell_{ok}^E$ , often much smaller than what could infer directly from the triangular relation of Eq. 16, e.g.,  $\ell_{ok}^E < \tilde{\ell}^B + \ell^\Phi$ . Indeed, one could presume that taking the turnaround point of the contour for a given accuracy could be the best choice here. In fact, such a choice would merely minimize the sum of both bandwidths, (or some monotonic function of each of them) for the given accuracy, what may or may not be relevant for a specific case at hand. Instead we would rather select the bandwidths to minimize explicitly actual computational cost of whatever code we plan on using. We present specialized considerations of this sort in the next Section.

On a more general level, we find that the standard rule of thumb, interpreting the effects lensing as a convolution of the unlensed CMB signal with a relatively narrow,  $\Delta\ell \sim 500$ , convolution kernel due to the lensing potential applies only for  $T$  and  $E$  signals and even in these cases only to small and intermediate values of  $\tilde{\ell}^{T,E} \lesssim 2000$  and only as long as the computation precision on order of  $\sim 1\%$  is sufficient. For higher values of the lensed spectrum multipoles or higher levels of the desired accuracy in the case of  $T$  and  $E$  and for all multipoles of the  $B$ -polarization signal the required bandwidths of both the respective, unlensed CMB signal and the gravitational potential are more comparable and indeed the latter bandwidth is often found to be larger.

We note that the analysis of this sort is somewhat more prone to problems in the case of  $TE$  power spectrum since the lensing kernels  $\mathcal{K}_{\tilde{\ell}^{TE}}(\ell^{TE}, \ell^\Phi)$  are not always positive as they contain the products of two different Wigner 3j coefficients and the  $TE$  power spectra, which may be non-positive, rendering the corresponding accuracy function not strictly monotonic. Hereafter, we will exclude this spectrum from our analysis, noting that any band limits prescriptions derived for  $T$  and  $E$  will also apply directly in the case of  $TE$ .

## 4. Numerical analysis

In this Section, we present numerical results concerning simulations of the lensed polarized spectra. Our goal here is two-fold. First, we study numerical consistency of the pixel-domain approach to simulating the lensing effect. Then, we demonstrate how the consideration from the previous section can be used to optimize numerical calculations involved in such simulations.

We start this Section from introducing a new implementation of the pixel-domain algorithm, which we refer to as **lenS<sup>2</sup>HAT**.

### 4.1. lenS<sup>2</sup>HAT

**lenS<sup>2</sup>HAT** is a simple implementation of the pixel-domain algorithm for simulating effects of lensing on the CMB anisotropies. The hallmark of the code is algorithmic simplicity and robustness, with the performance being assured by an efficient, distributed memory parallelization, making it a suitable choice for



massively parallel supercomputers, while its implementation follows the blueprint proposed originally by Lewis (2005) and summarised in Sect. 2.2.1. The main features of the code are as follows.

**Grids.** The code can produce lensed maps in a number of pixelization popular in the cosmological community, however internally it uses grids based on Equidistant Cylindrical Projection (ECP) where grid points, or pixel centers, are arranged in a number of equidistant iso-latitudinal rings, with points along each ring being again equidistant. Such a pixelization supports a perfect quadrature for band limited functions, what in the context of this work permits minimizing undesirable leakages, typically plaguing codes of this type. It can be shown, Driscoll & Healy (1994), that an ECP grid made of  $2L$  isolatitudinal rings, each with  $2L$  points and a weight, as given by,

$$w^j = \frac{2\pi}{L^2} \sin(\theta_j) \sum_{\ell=0}^{L-1} \frac{\sin((2\ell+1)\theta_j)}{2\ell+1}, \quad \theta_j = \frac{\pi}{2L} j, \quad (21)$$

is required and sufficient to ensure a perfect quadrature for any function with a band not larger than  $L$ .

**Interpolation.** For the interpolation, the code employs the nearest grid point (NGP) assignment, e.g., we assign to every deflected direction a value of the sky signal computed at the nearest grid point which is defined by the centers of the pixel of assumed pixelization scheme and therefore the respective sky signal values are calculable at the fast spherical harmonic speed. The NGP assignment is extremely quick and simple however it requires the computations to be performed at a very high resolution to ensure that the results are reliable. The sufficient resolution required for this will in general depend on the intrinsic sky signal, prior to the lensing procedure, as well as the resolution of the final maps to be produced as it is discussed in Sect. 4.2. As it is already clear from our theoretical discussion earlier in a typical case these are expected to be very high and the computations involved in the problem may quickly become prohibitively expensive.

**Spherical harmonic transforms.** To sidestep the problem of computing spherical harmonic transforms with a huge number of grid points and very large band limit,  $\text{lenS}^2\text{HAT}$  resorts to parallel computers and massively parallel numerical applications. With those becoming quickly more ubiquitous and affordable such a solution is becoming progressively more attractive.

Parallelization of the fast spherical harmonic transforms is difficult due to the character of the input and output objects and the involved computations, where a calculation of each output datum requires knowledge of, and access to, all input data. This is clearly not straightforward to achieve without extensive data redundancy, as done e.g., in  $\text{LensPix}$  or parallel routines of  $\text{HEALpix}$ , or complex data exchanges between the CPUs involved in the computation. To avoid such problems in our implementation we use the publicly available Scalable Spherical Harmonic Transform ( $\text{S}^2\text{HAT}$ ) library<sup>6</sup>. This library provides a set of routines designed to perform Fourier analysis of arbitrary spin fields on the sphere on distributed memory architectures (though it has an efficient performance even when working in the serial case). It has a nearly perfect memory scalability obtained

via a memory distribution of all main pixel and harmonic domain objects (i.e., maps and harmonic coefficients), and ensure very good load balance from the memory and calculation points of view. It is a very flexible tool allowing simultaneous and multi-map analysis of any iso-latitude pixelization, symmetric with respect to the equator, with pixels equally-distributed in the azimuthal angle, and provides support for a number of pixelization schemes including the above mentioned ECP. See, Szydlarski et al. (2011) for more details. The core of the library is written in F90 with a C interface and it uses the Message Passing Interface (MPI) to institute distributed memory communication, what ensures its portability. The latest release of the library also includes routines suitable for General Purpose Graphic Processing Units (GP-GPUs) coded in CUDA (Hupca et al. 2012; Szydlarski et al. 2011; Fabbian et al. 2012).

We emphasize that if a sufficient resolution can be indeed attained the approach implemented here can produce results with essentially arbitrary precision. In the following we will demonstrate that it is indeed the case for the described code.

## 4.2. Code parameters

### 4.2.1. Overview


In this Section we review how we fix the essential code parameters. These are introduced here with an emphasis on the causal relations between them, while a detailed description of the procedures used to assign specific values to them, is presented in the following sections.

1. We start from defining a target value in terms of the maximal value of the harmonic mode,  $\tilde{\ell}_{req}^Y$ , we want to recover and its desired precision,  $\varepsilon$ . We use then the reasoning from Sect. 3.2 to translate this requirement into corresponding bandwidths,  $\ell^X$  and  $\ell^\Phi$ , of the relevant unlensed signals,  $X$  and  $\Phi$ . These assure that the precision of all modes of the lensed signal up to  $\tilde{\ell}_{req}^Y$  be not worse than  $\varepsilon$ , barring any unaccounted-for, numerical inaccuracies. The values of  $\ell^X$  and  $\ell^\Phi$  are used to estimate the band limit of the output, lensed map,  $\tilde{\ell}_{out}^Y$ .
2. We then simulate two unlensed maps,  $\mathbf{m}^X$  and  $\mathbf{m}^\Phi$ , of the signal  $X$  and potential field,  $\Phi$ , with their band limits set to  $\ell^X$  and  $\ell^\Phi$  as estimated earlier. The number of pixels of the displacement map,  $\mathbf{m}^\Phi$ , is equal to that in the output map of the lensed signal, and for the ECP grid, equal therefore to  $N_{pix}^\Phi = 4 \tilde{\ell}_{out}^Y{}^2$ . The number of pixels in the  $X$ -signal map,  $\mathbf{m}^X$  is then given by  $N_{pix}^X = 4 \kappa^2 \ell^X{}^2$ , where  $\kappa$  is an overpixelization factor introduced already in Sect. 2.2.1 and discussed in detail below, Sect. 4.2.4. For simplicity, we assume that the grid for which the unlensed field  $X$  is computed is a subgrid of the grid used for  $\Phi$ .
3. The reassignment procedure (step 5 of the algorithm, Sect. 2.2.1) is then straightforwardly performed leading to the map containing power potentially up to  $\tilde{\ell}_{out}^Y$ , which may need to be filtered down to the band limit of  $\tilde{\ell}_{req}^Y$ , as initially required.

### 4.2.2. Intrinsic bandwidths

We employ the procedure described earlier in this work in Sect. 3.2 to set the intrinsic band limits. Rather than to use generic predictions we aim at optimizing their values to ensure the lowest possible computational overhead. To do so we need to

<sup>6</sup> [http://www.apc.univ-paris7.fr/APC\\_CS/Recherche/Adami/MIDAS09/software/s2hat/s2hat.html](http://www.apc.univ-paris7.fr/APC_CS/Recherche/Adami/MIDAS09/software/s2hat/s2hat.html)



**Fig. 6.** Numerical cost gain in using optimized set of  $\ell^\Phi, \ell^E$  parameters with respect to the assumption of  $\ell^\Phi = \ell^E$  as a function of accuracy computed assumption for several multipoles. An oversampling factor of  $\kappa = 8$  is assumed to compute the cost function

provide a model of the cost of numerical calculations involved in  $\text{lenS}^2\text{HAT}$ . This is dominated by large spherical harmonic transforms, one to calculate the map of  $\Phi$  and the other that of  $X$ . Given the parameters introduced above and the fact that the total cost of a spherical harmonic transform is proportional to  $N_{\text{pix}} \ell_{\text{max}}$  we therefore obtain,

$$C \equiv C(\ell^\Phi, \ell^X) \propto 2 N_{\text{pix}}^\Phi \ell^\Phi + n_{\text{stokes}} N_{\text{pix}}^X \ell^X \quad (22)$$

$$= 8 \tilde{\ell}_{\text{out}}^2 + 4 n_{\text{stokes}} \kappa^2 \ell^X \quad (23)$$

$$= 8 \eta^2 (\ell^\Phi + \ell^X)^2 \ell^\Phi + 4 n_{\text{stokes}} \kappa^2 \ell^X. \quad (24)$$

Here  $n_{\text{stokes}}$  stands for a number of signals maps of which we want to produce and is equal 1 –  $T$ -only, 2 –  $E$  and  $B$ , or 3 –  $T$ ,  $E$ , and  $B$ , while for the field  $\Phi$  the pre-factor is fixed and equal to 2 reflecting the number of components of a vector field on the sphere. In deriving the last equation above we have assumed that  $\tilde{\ell}_{\text{out}} = \eta (\ell^\Phi + \ell^X)$  as is justified below, as are the values which should be adopted for  $\eta$  and  $\kappa$ . The expression above includes neither the cost of the interpolation nor reshuffling but as both these operations are proportional to  $N_{\text{pix}}$  their cost is negligible with respect to the cost of the transforms.

Solving for the optimized values of the bandwidths, which simultaneously ensure desired precision at  $\tilde{\ell}_{\text{req}}^Y$ , as expressed by  $\varepsilon$  involves now minimizing the cost function in Eq. 24, with a constrain  $A_{\tilde{\ell}_{\text{req}}}^Y(\ell^\Phi, \ell^X) = \varepsilon$ , on Eqs. 17 and 20. We implemented this operation as a two step procedure. Given that the cost function is a smooth function, we first defined a grid of levels on it and then computed the maximum of the accuracy function along the cost contours. The location of this maximum represents then our best set of bandwidth parameters. As a result of optimization procedure we end up having a set of values  $(\ell^\Phi, \ell^X)$  fixing the best accuracy which can be reached for a given numerical cost provided by that specific choice of parameters. In general those are not equal and there is no general argument which prevents to keep the band limit of CMB equal to the one in  $\Phi$  as it is assumed in the current available methods (in the following we will refer to this assumption as the *diagonal assumption*). The result of the optimization procedure is driven by the shape of the lensing kernels and it is more fruitful when simulating the CMB modes spectra at very high multipoles, in particular for  $B$  modes which have a broader window functions and are more demanding in terms of bandwidth requirements. In the latter case in fact we need to include progressively more power from lensing potential with respect to CMB to properly take into account the contribution of the tail. The optimized parameters therefore deviate more from the assumption  $\ell^\Phi = \ell^X$  leading to a cost increased only along one direction and an overall decrease of the

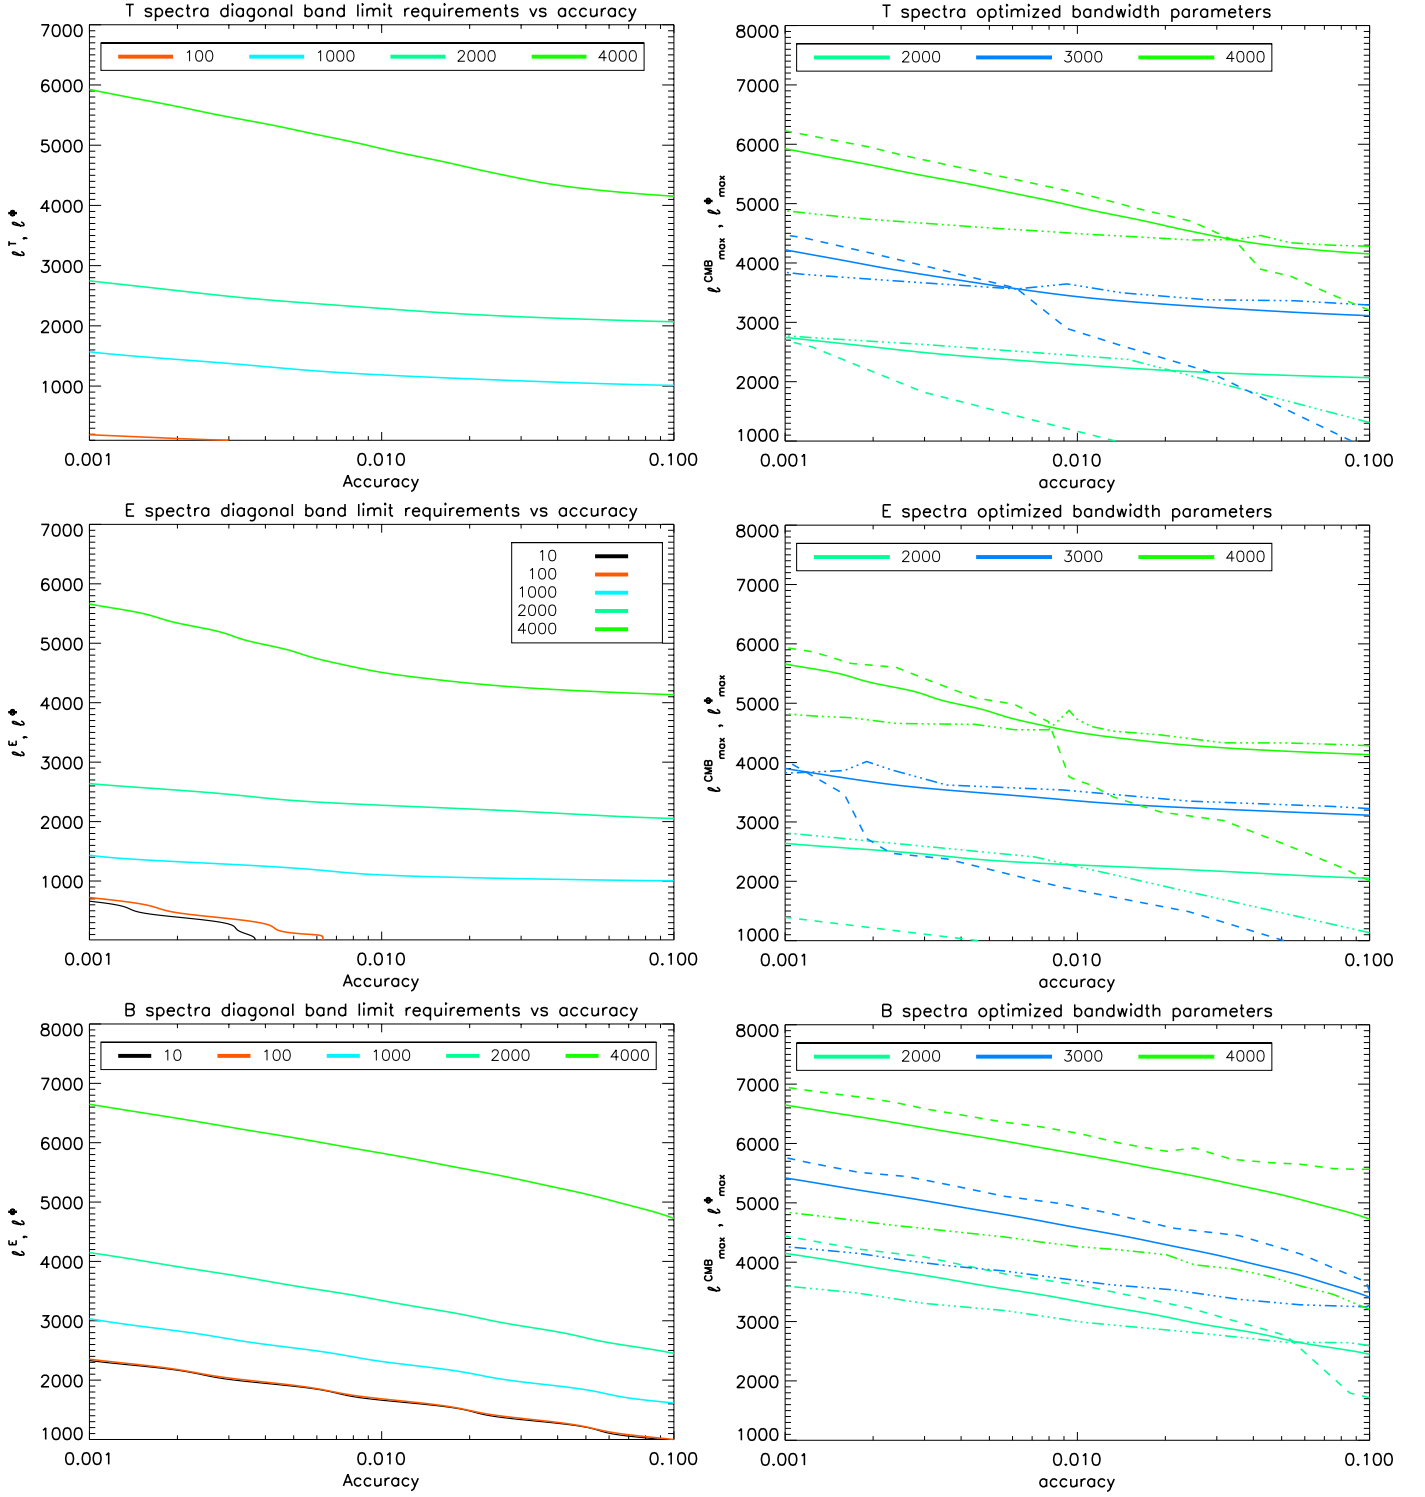
cost for a given accuracy. For the regimes we considered, the reshuffling of power from smaller angular scales  $\ell$  to larger angular scales leads to a significant contribution if it is correlated with very small scale power in lensing potential i.e. if the both peak and the low  $\ell$  tail of that mode is properly simulated. We also note that the optimized bandwidth parameters as a function of accuracy present a crossing point whose location is set by the location and broadness of the peak contribution of the lensing kernels (see. fig. 3, 4). Around unlensed mode equals to the lensed one in exam, the peak of the kernels, especially for intermediate scales, is broader in the CMB direction rather than in the  $\Phi$  direction and is therefore more convenient to include more power in that direction. The oversampling factor, if sufficiently high, penalizes the extension in this direction and can shift the position of the crossing point.

The overall procedure allows to gain for lensed  $B$  multipoles close to 4000, a factor of nearly 40% in terms of runtime in a range of accuracy of interest if an highly oversampling factor for  $N_{\text{pix}}^{\text{CMB}}$  is required. For temperature and  $E$  modes polarization, where less extra power is required in  $\Phi$  to obtain an accurate result, the gain can be quantified in nearly 20% - 30%. We report in figure 7 the results of bandwidth parameters optimization for different lensed multipoles of  $T$ ,  $E$ ,  $B$  spectra and the correspondent parameters providing the same accuracy if the diagonal assumption is followed while in figure 6 we report the average value of runtime gains computed as the average fractional difference on cost function for a range of accuracy  $10^{-3}, 10^2$ . It is interest to note that both the optimized and diagonal band limit requirements to obtain a lensed spectra up to a multipole  $\ell$  with a given accuracy, guarantee that the same accuracy level is reached for lensed multipoles  $\ell' < \ell$  especially if  $B$  modes bandwidth requirements are followed.

#### 4.2.3. Lensed map band limit

For the resolution of the final map, we note that in an absence of numerical effects, such as those due to the pixelization and interpolation, the lensing procedure would be described by Eq. 12 and the bandwidth of the lensed map would be simply given by  $\ell^X + \ell^\Phi$ . In the presence of the numerical effects, the output map will have the effective bandwidth typically higher than that, what will typically lead to some power aliasing at the high- $\ell$  end if this theoretical band limit is imposed. We find this to be indeed the case in our numerical calculations. However, we also find that, once the overpixelization factor is set correctly, the aliasing is localized to at most 25% of the bandwidth and therefore easy to deal with in post-processing, e.g., step 6 of the algorithm outline






**Fig. 7.** Left column: Examples of band limit requirements for unlensed CMB and lensing potential as a function of obtainable accuracy, assuming those to be both equal, on several multipoles of lensed T, E, B spectra. Right column: summary of optimized choice of bandwidth parameters for CMB (dot dashed) and  $\Phi$  field (dashed) with respect to the cost function of the algorithm. As defined in eq. 24. The diagonal bandwidth parameters are shown in solid line as a comparison.

in Sect. 2.2.1. Consequently in our numerical simulations we use  $\tilde{\ell}_{out}^Y = \eta(\ell^X + \ell^\Phi)$ , with  $\eta = 1.25$  as the band limit.

It is important to emphasize that NGP is one of the sources of the aliasing, as it does not preserve the bandwidth of the interpolated function, as do not some of the other procedures proposed ad hoc in this context. Clearly, an interpolation preserving the

function bandwidth would be a clear improvement for this kind of algorithms, if it comes without prohibitive numerical cost. We leave such an investigation to future work.




**Fig. 8.** Comparison between the E modes power spectra of input displacement field (black) and the displacement field after NGP assignment for several oversampling factor  $\kappa$ . The input displacement is computed on an ECP grid with a number of pixel  $N_{pix} = 16384^2$  while the discretized one is the result of an NGP assignment on a grid of  $\kappa^2 N_{pix}$ . With progressively higher resolution the extra power due to discretization becomes negligible and the two spectra become almost indistinguishable. The discretization induced errors power spectra is shown in dotted line for reference; both E and B modes of discretization error have the same power spectra.

#### 4.2.4. Overpixelization factor

As was explained already earlier our interpolation procedure consists in two steps: an overpixelization followed by an NGP assignment. The overpixelization involves producing maps with sky signal sampled at significantly higher rate than what necessary given the signal's band limit. In the case of the ECP grids used internally by `lenS2HAT` this is implemented by using  $\kappa$  as many points in each of the two,  $\phi$  and  $\theta$  directions. The remaining problem is then fixing the appropriate value of  $\kappa$ . To do so we first observe that for the overpixelized grid, the NGP assignment can be seen either as approximating the true value of the sky signal, which needs to be calculated in one of the displaced direction, computed in turn precisely – this is the standard perspective and the only available if a more sophisticated interpolation scheme is applied – or as approximating the displaced direction by one pointing towards the nearest grid point, while assigning to it a correct sky value. This second viewpoint provides as with an independent check whether our overpixelized grid is sufficiently dense. The involved procedure involves first calculating the approximate displacement field and calculating its power spectrum, which is then compared against the input power spectrum for the gravitational potential,  $\Phi$ . We note that the approximation used here can in general generate a non-zero curl and therefore there will be two non-vanishing spectra of the approximate displacement field, corresponding to its  $E$  (gradient) and  $B$  (curl) components. We then require that the recovered  $B$  spectrum is significantly smaller than the  $E$  one, and that both the recovered  $E$  spectrum and the input one agree sufficiently well up to the angular scales, which are of interest given the  $\ell$ -range of the lensed spectrum we are after and its precision. These latter two are turned into the  $\ell$ -range requirement using one of the 2-dimensional kernels.

Examples of such comparisons are shown in Fig. 8 for a number of values of the oversampling factor ranging from  $\kappa = 1$  up to 8. We see that for the latter value the approximate  $E$  spectrum is consistent over the entire shown range of  $\ell$  values and the recovered  $B$  is there significantly smaller. We therefore will keep on using this value in the runs discussed later in this work, even if, as noted in the next Section,  $\kappa = 4$  could be sufficient at least for  $\tilde{\ell}^B \lesssim 2000$ . We also point out that, as it could have been expected, the departures of the recovered  $E$  spectrum for the displacement from the input one, are consistent with the presence of the non-zero  $B$ -type mode in the approximated displacement field with an amplitude comparable to that of its  $E$ -mode spectrum, rendering therefore our two criteria in fact redundant. In addition, if only  $T$  and  $E$  CMB spectra are of interest then  $\kappa \approx 2$  is already sufficient to obtain accurate result on the scales of interest given that the long-tails of the displacement spectrum do not matter much in these cases.

For completeness, in Fig. 9 we show the relevant CMB  $B$ -spectra computed with the same values  $\kappa$  as shown in Fig. 8 and aiming at a high precision reconstruction for  $\tilde{\ell}^B \leq 2000$  demonstrating that both overpixelization rates, as inferred above, ensure a satisfactory recovery of this spectrum in the targeted range of  $\ell$ . We provide further details about this Figure in Sect. 4.3.2.



**Fig. 9.** Lensed  $B$  mode spectra computed for different value of over sampling factor compared to the lensed spectrum obtained with the analytical Boltzmann code CAMB (red dashed).

### 4.3. Validation and tests

#### 4.3.1. Simulated kernels

As a first step of validation of our code we investigated whether its results agree with the prediction of the semi-analytical approach used to model convolution in the harmonic domain. We focus here on numerically feasible studies of the one-dimensional kernels, as defined in Eq. 13. For this purpose we assume that the unlensed CMB signal, i.e.,  $E$ -mode polarization in the case of the lensed  $B$ -spectra, contains power only in a single harmonic mode,  $\ell_0$  i.e.,  $C_\ell^{EE} \propto \delta_\ell^{\text{Kronecker}} \ell_0$  and compute using `lenS2HAT` the resulting lensed  $B$  mode spectra for several values of  $\ell_0$  and compare them with the analytic results obtained for the same multipole and displayed in Fig. 2. The results of this calculation are shown in Fig. 10, where we see that in a range where

the analytic model is more reliable the agreement between the two curves is excellent if a sufficient resolution for the unlensed grid is used, and the simulations seem to confirm the analytic results. On the other hand, in the region where the analytic approximation we used is not anymore accurate, since the CMB signal amplitude and the one of its gradient have comparable magnitude and thus truncation in the series expansion introduce a non negligible error, the discrepancy between our analytical model and the simulated one dimensional spectra becomes more evident. Such an approximation tends to over estimate the contribution of each single mode to its neighboring angular scales of a factor of nearly 50% with respect to simulated kernels and to slightly under estimate the contribution to each mode to the tail of the kernels, i.e. to the multipoles farther to itself.

We emphasize that these results not only do not undermine the validity of the band limits requirements we have derived in the previous sections, but in the contrary, they validate them by demonstrating that the approximate and actual kernels are qualitatively rather similar, what is sufficient for our purposes as the definition of the accuracy function is less sensitive to errors in the amplitude of contribution of a single pairs of modes. Moreover an over estimation of contribution of local high  $\ell$  power to the total spectrum through the analytical model, is translated in more conservative requirements to achieve a given accuracy and the recipe provided in previous section should not lead to any underestimation of accuracy for our given input parameters.

#### 4.3.2. Simulated spectra

Another batch of performed tests involves a comparison of the spectra obtained from `lenS2HAT` and those derived with Boltzmann codes such as `CAMB` or `CLASS`. In particular, in Fig. 9, the black solid line shows an example of the result obtained for a simulation of lensed  $B$  modes designed to reach an accuracy up to 0.1% at  $\ell^B \lesssim 2000$ . We note that, as no band limit optimization is performed, and therefore it is assumed that  $\ell^\Phi = \ell^E = \ell_{max}$ , the latter value has to be at least  $\ell_{max} = 4000$ , Fig. 4. The lensing convolution of signals having such a band limit leads to polarized maps with power up to  $2\ell_{max}$  and thus to avoid any aliasing we would need a grid for the lensed sky and the displacement field having at least  $N_\theta \approx 2\ell_{max}$  rings with  $N_\phi \approx 2\ell_{max}$  pixels per ring, i.e.  $N_{pix} \approx 16384^2$ , where we have rounded the number of rings and pixels per rings to a power of 2. Once the band limit of the signals and the respective grid for the lensed sky is set we still need to define the overpixelization rate as required by our interpolation. As noted in the previous Section, there seem to be a general argument based on the discretized displacement spectra, which points towards  $\kappa = 8$  as a sufficient value. As calculating the overpixelized map, albeit with a fixed band limit, is the most time consuming part of the code, there may be on occasions an interest in tuning  $\kappa$  as small as only possible. In this context we find, as illustrated in Figs. 8 and 9, that if the extra power introduced by discretization of displacement field does not exceed 10% of the power in the non-discretized displacement field on scales  $\ell \approx 1.5\ell^B$ , then an oversampling factor of 4 is sufficient to render a power spectrum on scales  $\ell \lesssim \ell^B$  with an accuracy as determined by the assumed bandwidth. The factor 4 should be however treated as a lower bound and used with care, as there will be typically significant amount of extra power in the  $B$ -mode spectrum for  $\ell \gtrsim \ell^B$ , which may need to be efficiently filtered out before the respective map could be further used. In contrast, if the extra

power found in the discretized displacement does not exceed 10% of the original power for all angular scales up to  $\ell_{max}^\Phi$  then no overshooting takes place and results remain highly accurate also beyond the scale of interest  $\ell^B$ .

In Fig. 11 we present the spectra computed for both polarized components  $E$  and  $B$  as well as the displacement field,  $\Phi$ , for a run aiming at a high precision, better than to 0.1%, recovery of these signals in a band up to  $\tilde{\ell}^X \lesssim 5500$ . For this runs we assumed the value of the required bandlimits to be  $\ell_{max}^E = \ell_{max}^\Phi = 8000$ . Those have been extrapolated from Fig. 7, where to obtain 0.1% accuracy on  $B$  modes on similar angular scales (e.g.  $\tilde{\ell}^X = 4000$ ) we need to include power up to  $\ell_{max} \approx \tilde{\ell}^X + 2500$ . Following the same prescription given for the previously detailed case of Fig. 9, we set the resolution of unlensed sky and displacement field to  $N_{pix} = 32768^2$  while, to ensure the highest possible reliability of the result, we pushed the oversampling factor to 16. The errors discretization errors introduced by this setup are found to stay under the 1% level on all the angular scales involved in the calculation and no significant overshooting is shown (see. Fig 11). Though the band limit and resolution involved may look exaggerated from the practical point of view they simultaneously demonstrate the capability of the numerical code, while illustrating our conclusions concerning the precision of such calculations.


In general, we find that a simple algorithm as proposed in `lenS2HAT` is capable of simulating effects of lensing on CMB over the range of angular scales of interest for current and foreseeable experimental efforts. Moreover, if used properly, it does so with the accuracy, which on very small scales is limited rather by the precision of the input power spectrum of unlensed CMB than the employed numerical algorithm.

#### 4.4. Convergence tests

In order to investigate the precision and reliability of our approach it is interesting to investigate the numerical convergence of the results without relying on a direct comparison to an external Boltzmann code. Since several experiments in the future will be able to target non-gaussianities in CMB polarization i.e. statistical moments beyond the power spectrum, we decided to study the convergence of the results not only on the power spectrum level but also in the real domain i.e. on the map level.

#### 4.5. Power spectrum convergence

We first investigate the convergence of the power spectrum up to a given scale of interest  $\tilde{\ell}^X$  as a function of the bandlimits. This procedure would allow us to show simultaneously the precision of our code and also indirectly prove the validity of the bandwidth requirements given in Sect. 4.1. For this purpose we assumed the bandlimits to be equal for CMB and  $\Phi$  field and we then fixed resolution of the grid following the prescriptions of Sect. 4.3.2 assuming  $\kappa = 8$ . We simulated a full CMB maps with all the 3 Stokes parameters  $T$ ,  $Q$ , and  $U$  and then compute the precision on the value of the multipole  $\tilde{\ell}^X$  of interest recovered from the simulation as the fractional difference between the values obtained from two subsequent values of  $\ell_{max}$  considered. For this specific tests we made sure that the random realization of the harmonic coefficients used in the simulation is the same when changing the value of the bandlimit from  $\ell_{max}$  to a value  $\ell'_{max}$  for  $\ell \leq \ell_{max}$ . We report the result of the numerical convergence for



**Fig. 10.** Comparison between simulated, solid black lines, and analytical, solid red lines, 1-dimensional B modes kernels,  $\mathcal{H}_{\ell^E}(\tilde{\ell}^B)$ , shown as a function of  $\tilde{\ell}^B$ , and for the unlensed CMB  $E$  power contained initially only in a single mode,  $\ell^E = 500, 1000, 4000, 5000$ . The analytic expression, Eq. (12), reproduces the kernels obtained from simulations nearly perfectly for  $\ell^B \lesssim 2000$ , e.g., in the regime where the gradient approximation is more reliable (left panels). On smaller angular scales (right) the agreement is worse, though still qualitatively very good and the observed discrepancies pose no for the semi-analytic considerations presented in Sect. 3.

$\tilde{\ell}^X = 2000$  in Table 4.5. We note that the results are in agreement with the analytic calculation of Sect. 4.1 where we saw that extending the band limit has no visible effect in the recovered results on the scale of interest if a proper amount of power has already been convolved. We note that, as expected, a significant fraction of  $E$  modes power is converted into B-modes for angular scales  $\ell^E \in [3000, 4000]$  but no significant improvement is seen if power beyond  $\ell^E = 4000$  is included. The test case for  $\tilde{\ell}^X = 4000$ , i.e. in the regime where the gradient approximation is expected to be less accurate was also been performed and the results are summarized in table 4.5. The B modes accuracies are consistent with the one derived in Sect. 4.3.2 except for the last set of bandwidth parameters, where the fractional difference between simulated spectra seems to saturate to a level of 0.1%. The latter may be explained to a small residual aliasing effect due to an underestimated oversampling parameter since an increasing of  $\kappa$  would lead to an improvement on the result  $\lesssim 0.1\%$

$\tilde{\ell}^X$	2000	3000	4000	6000
TT	43%	0.04%	0.02%	0.003%
EE	31%	0.01%	0.01%	0.005%
BB	35%	3%	0.02%	0.004%
TE	32%	0.04%	0.01%	0.002%

**Table 1.** Numerical convergence of simulated CMB power spectra multipole  $\tilde{\ell}^X = 2000$  as a function of assumed bandlimit of unlensed CMB and  $\Phi$  field used in the simulation. The  $i$ -th column show the fractional difference between results computed using two subsequent values of  $\ell_{max,i}, \ell_{max,i+1}$  in the set  $\ell_{max} \in \{2000, 3000, 4000, 6000, 8000\}$

#### 4.6. Map convergence

After having shown the convergence on the power spectrum level, which provides information on the overall variance of the simulated maps, we investigated if the convergence of our numerical result is also realized in the real domain. For this pur-

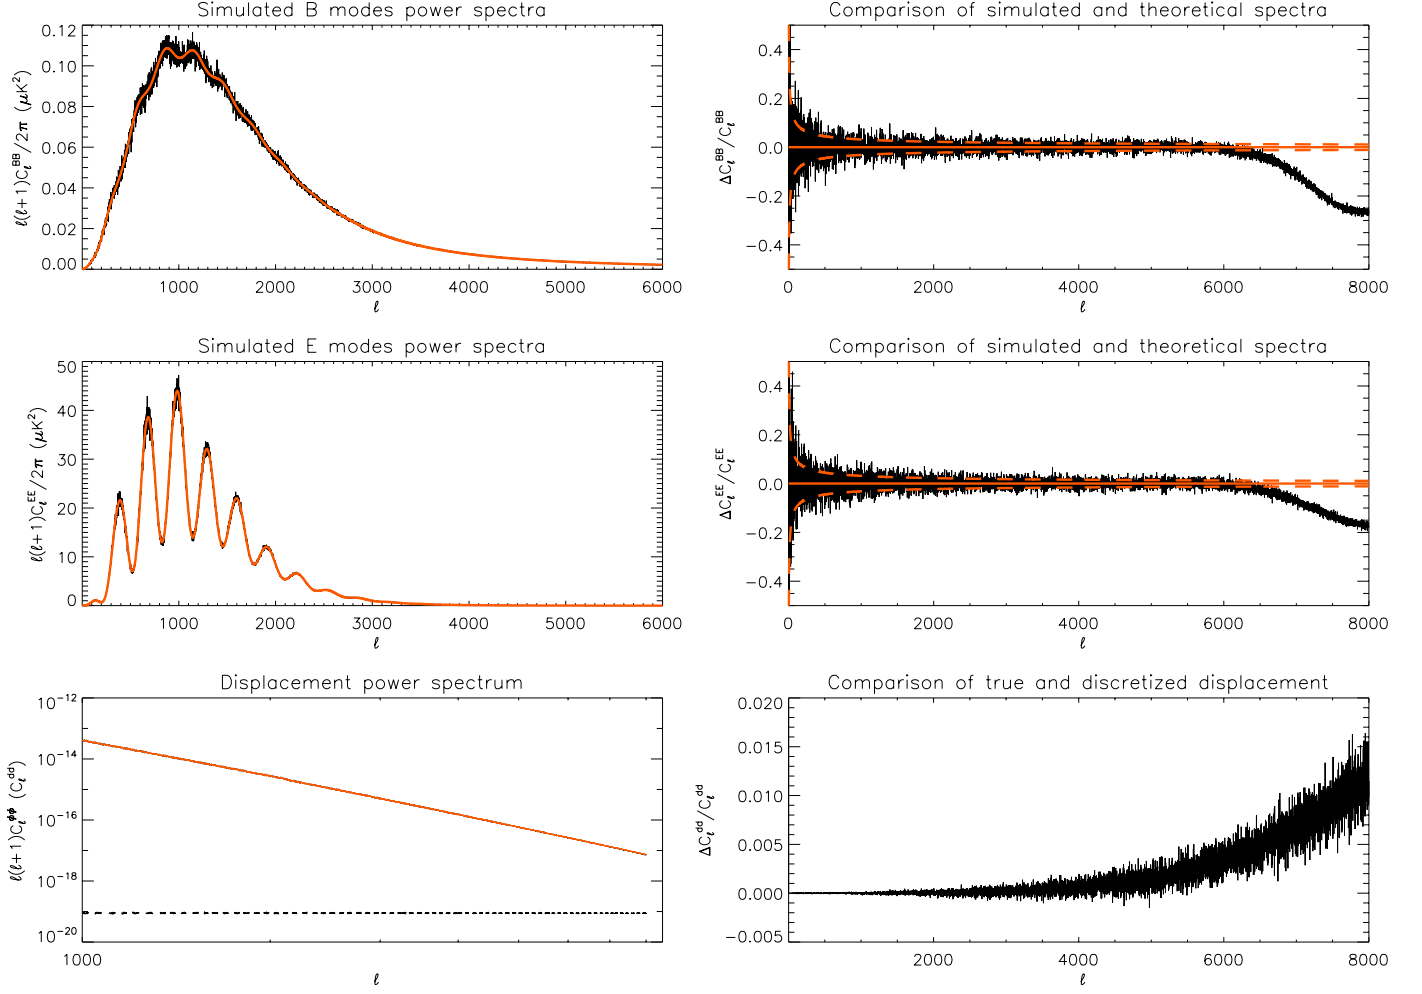
$\tilde{\ell}^X$	4000	5000	6000	7000
TT	31%	0.2%	0.09%	0.1%
EE	32%	0.2%	0.05%	0.03%
BB	7%	4.6%	0.1%	0.09%
TE	21%	0.7%	0.2%	0.2%

**Table 2.** Numerical convergence of simulated CMB power spectra multipole  $\tilde{\ell}^X = 4000$  as a function of assumed bandlimit of unlensed CMB and  $\Phi$  field used in the simulation. The  $i$ -th column show the fractional difference between results computed using two subsequent values of  $\ell_{max,i}, \ell_{max,i+1}$  in the set  $\ell_{max} \in \{4000, 5000, 6000, 7000, 8000\}$


pose we first define an error map obtained as a difference of two maps computed assuming two different bandlimits  $\ell_{max,1}, \ell_{max,2}$  rescaled by the rms value of one of the two maps, i.e.,

$$\mathcal{E}_{\ell_{max,1}, \ell_{max,2}}^X = \frac{\mathbf{m}_{\ell_{max,2}}^X - \mathbf{m}_{\ell_{max,1}}^X}{\sqrt{\text{Var}(\mathbf{m}_{\ell_{max,2}}^X)}} \quad X \in \{T, Q, U\} \quad (25)$$

where  $\mathbf{m}_{\ell_{max,1}}^X$  is a simulated map of the field  $X$  obtained assuming  $\ell_{max,1}$  as the bandlimit. Once we have derived the harmonic coefficients with the procedure outlined in the previous section, we filter out all the modes on angular scales  $\ell \geq \tilde{\ell}^X$  and we resample the signal on grid which properly samples the signal up to multipole  $\tilde{\ell}^X$ . In order to take advantages of the HEALPIX visualization tools we use for this purpose an HEALPIX grid having  $n_{side} = 1024$  (2048) for  $\tilde{\ell}^X = 2000$  (4000). After having resampled the harmonic coefficients we then compute the power spectrum of the error field  $\mathcal{E}_{\ell_{max,1}, \ell_{max,2}}^X$ , which demonstrates the precision obtained on the map level. In Fig. 13 we report the result of this analysis for the test case  $\tilde{\ell}^{T,E,B} = 2000$ . The error field power spectrum is found to be very similar to a white noise spectra with r.m.s below 1%. If the power is properly resolved an accuracy equivalent to 0.1% on the map level can be indeed been obtained while for polarization the error slightly increases to the 0.3% level (see Fig. 14). We note however that this test case does not include the effect of any realistic experiment




**Fig. 11.** Example of a single realization of a high accuracy and high precision simulation of CMB polarization obtained with lenS<sup>2</sup>HAT. The run required an unprecedented resolution and bandwidth and was performed as a demonstration of the code capabilities. The bandwidth parameters have been set to obtain 0.1% accuracy of modes up to  $\ell \approx 5500$  for both  $B$  and  $E$  spectra. For the polarization signals, top and middle panels, right panels show the spectra, both recovered from the lenS<sup>2</sup>HAT run, black solid lines, and the theoretical one obtained from CAMB. Left panels show the relative difference between the two with the dashed lines outlining expected  $1\sigma$  uncertainty due to sampling variance. The bottom, left panel shows the  $E$  power spectrum of the recovered displacement *before* the NGP assignment, red solid line, and the  $E$  and  $B$  power spectrum of the displacement field *after* the NGP, black dashed lines, with the former  $E$  spectrum overlapping the latter nearly perfectly. The bottom, right panel shows then the relative discretization error.



**Fig. 12.** Performance benchmarks of the lenS<sup>2</sup>HAT code. From left to right, we show the memory per processor as a function of the number of processors (or MPI tasks), the total required memory and the number of floating point operations (FLOPs) as a function of the band limit and the walltime used per processor. In all cases the simulations have been done for three Stokes parameters and used a ECP grid of  $32768^2$  points,  $\ell_{max}^{CMB} = \ell_{max}^\Phi = 8000$  and oversampling factor  $\kappa = 8$ . In each panel the dashed lines show a theoretical scaling as expected in the ideal circumstances.



setup; in real life case, if instrumental noise has to be added to the simulated maps, the criterion for the convergence is set by the noise level on the pixel.



**Fig. 14.** T, E and B power spectrum of the error field obtained from maps simulated with different values of bandlimit parameters (colored lines). The T, E and B modes power spectra obtained of a gaussian random field on the sphere having 1% (dot-dashed) and 0.1% (dashed) variance are overplotted for comparison.

#### 4.7. Monte Carlo simulations

In order to test if our method produces any significant bias in the power spectrum we performed  $N_r = 100$  realization of  $T$ ,  $Q$ ,  $U$  lensed maps having requirements to obtain 0.1% accuracy up to  $\tilde{\ell}^B = 3000$  and investigate the statistical property of the power spectra averaged over  $N_r$ . The latter is in fact expected to be nearly gaussian distributed since the non gaussian correlation in the lensed power spectrum covariance induced by lensing itself are negligible for  $T$ ,  $E$  and  $TE$  spectra. For all the power spectra including the  $B$  field however, we expect the latter statement to be only partially correct since the covariance of such spectra is non-Gaussian, especially on small angular scale. Identifying the expected scatter of the averaged spectrum with the theoretical Gaussian sample variance tends therefore to underestimate the scatter itself.

For each couple of observed power spectra  $X$  and  $Y$  we defined the quantity

$$G_{\ell}^{XY} = \sqrt{\frac{(2\ell + 1) \cdot N_r}{[(\bar{C}_{\ell}^{XY})^2 + C_{\ell,th}^{XX} C_{\ell,th}^{YY}]}} (\bar{C}_{\ell}^{XY} - C_{\ell,th}^{XY}) \quad (26)$$

where  $\bar{C}_{\ell}^{XY}$  is the power spectrum averaged over  $N_r$  realizations. The ensemble of all the  $G_{\ell}^{XY}$  is expected to be Gaussian distributed with 0 mean and variance 1 and therefore we can perform a Kolmogorov-Smirnov test (denoted as KS in the following) if this hypothesis is verified. Starting from Eq. (26) and following (Basak et al. 2009) we can define a reduced  $\chi^2$  statistics in addition to the KS test as,

$$\tilde{\chi}_{XY}^2 = \sum_{\ell=2}^{\ell_{max}} \frac{|G_{\ell}^{XY}|^2}{\ell_{max} - 1} \quad (27)$$

We report in table 4.7 the results of both of those test expressed as the significance level probability of the null hypothesis. We

found that the method does not produce any significant bias also on TB and EB cross spectra; those have not been show in previous analysis but are themselves interesting since they are a sensitive test of any artificially induced correlation. We note moreover that the precision and accuracy of the result can be tested quite independently of analytical models devising a custom convergence procedure as explained in the previous section.

$\bar{C}_{\ell}^{XY}$	Significance $P_{KS}$	Significance $P_{\tilde{\chi}_{XY}^2}$
TT	0.19	0.92
EE	0.97	0.65
BB	0.79	0.14
TE	0.58	0.84
TB	0.20	0.34
EB	0.71	0.67

**Table 3.** Results of statistical tests on the recovered lensed power spectra averaged over 100 realizations. The significance level probability for the null hypothesis using a Kolmogorov-Smirnov test ( $P_{KS}$ ) and a reduced chi-square  $\tilde{\chi}^2$  statistics ( $P_{\tilde{\chi}_{XY}^2}$ ) show no bias on a statistical level

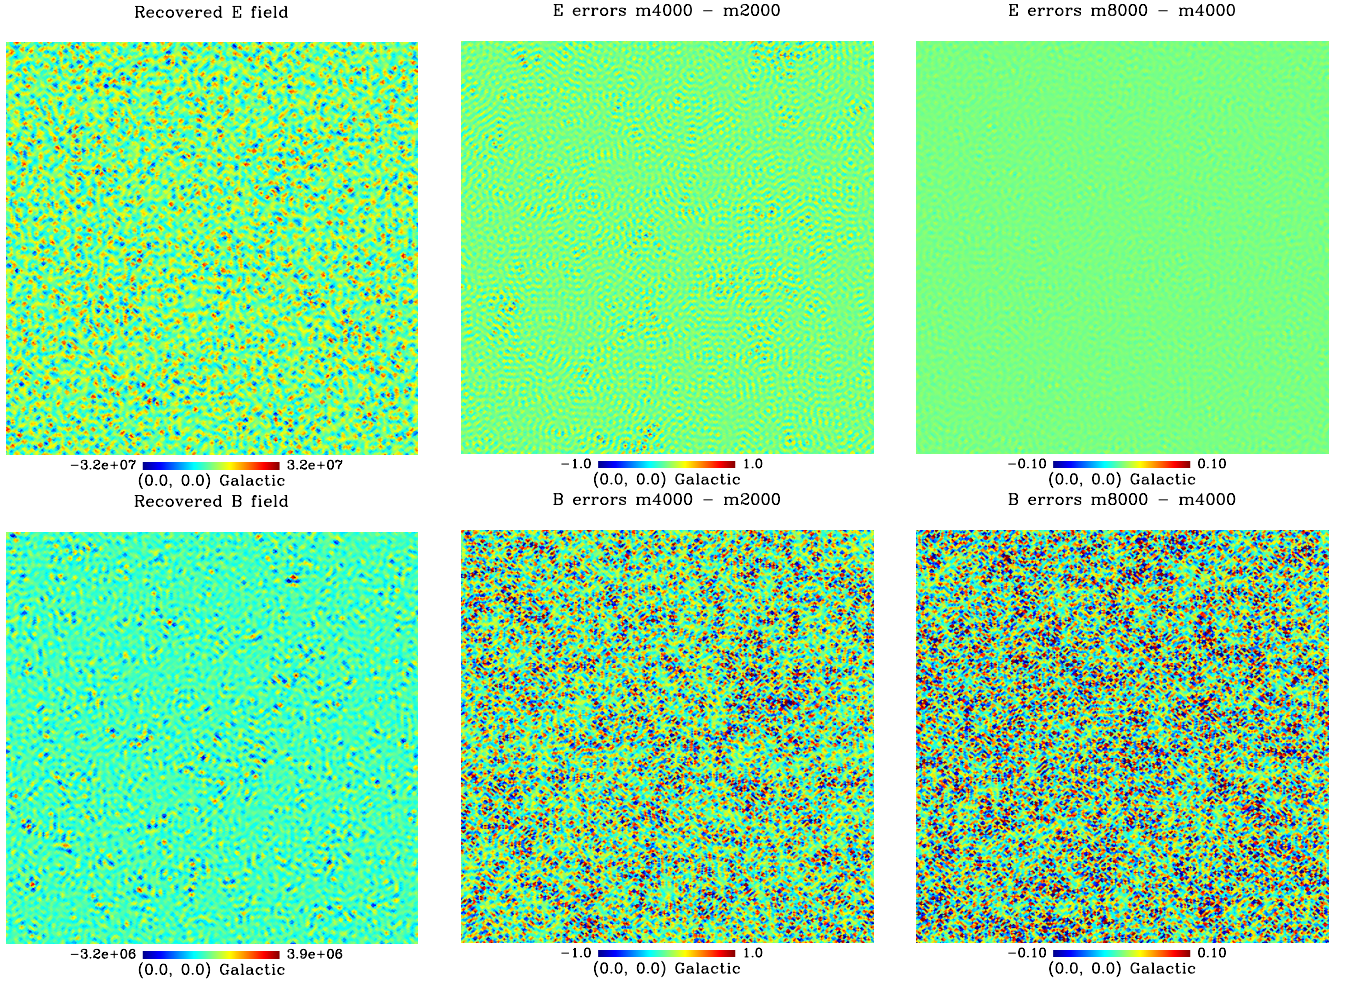
#### 4.8. Numerical performance and requirements

In this section we evaluate the weak scaling relations for both numerical cost and memory requirements of lenS<sup>2</sup>HAT, i.e. we run the code with the same parameters and test its scalability as a function of number of MPI processes used in the calculation. For this benchmark test we used,  $\ell_{max}^{E/B} = \ell_{max}^{\Phi} = 8000$  and a grid for lensed sky and displacement field of  $32768^2$  pixels.

The main data volume involved in the calculations is given by harmonic coefficients and maps which are evenly distributed between processors through the S<sup>2</sup>HAT library and their distribution is optimized for all the spherical harmonics transform steps involved. The remapping method itself only depends on structures which are also distributed between processors, allowing us to preserve the scalability features inherited by S<sup>2</sup>HAT library. The overall memory requirements per processors for a lenS<sup>2</sup>HAT run are of the order of  $\mathcal{O}(N_{pix}/n)$ , where  $N_{pix}$  is the total number of pixels of both lensed map and displacement field and  $n$  the number of MPI tasks (or processors) used for the simulation, which is assumed to be one for physical core available on our test architecture. The prefactor varies as a function of oversampling rate and on the type of simulation i.e. whether we want to simulate temperature maps only or also polarized maps or Q and U maps alone. For the first of the cases just outlined, the optimized factor is equal to  $(3 + \kappa^2)$  while for the other two cases the prefactor is equal to  $(4 + (1 + \kappa^2)n_{Stokes})$  where  $n_{Stokes}$  is the number of Stokes parameter to be simulated. We report in figure 12 the results of strong and weak scaling tests performed on the Cray XE6 Hopper cluster of the NERSC<sup>7</sup> supercomputing center using the Integrated Performance Monitoring library<sup>8</sup> (IPM). The discrepancy between our model and the actual memory resources requirements per processors are due to MPI buffers allocations for collective communications and duplications of auxiliary objects describing the properties of the pixelization and observed sky patch used in the simulations as required by the S<sup>2</sup>HAT library and remapping method. Those

<sup>7</sup> <https://www.nersc.gov>

<sup>8</sup> <http://ipm-hpc.sourceforge.net/>



**Fig. 13.** Maps of the recovered  $\chi^E$ , upper row, and  $\chi^B$ , lower, fields as defined in Zaldarriaga & Seljak (1997) (left column) and difference maps normalized by the input map rms,  $\mathcal{E}_{\ell_{max,1}, \ell_{max,2}}^{E/B}$ , Eq. (25), computed for  $(\ell_{max,1}, \ell_{max,2}) = (4000, 2000)$  (middle column) and  $(\ell_{max,1}, \ell_{max,2}) = (8000, 4000)$  (right). There is a factor of 10 difference between the colour scales of the middle and left column.

have a size  $O(11(1 + \kappa)\sqrt{N_{pix}})$  and accounts for nearly 25Mb of data duplicated on each processor. The memory overhead of communication part of the  $\text{lenS}^2\text{HAT}$  algorithm depends instead on the number of pixels in the local memory of each processors which is lensed on an area of the map stored on the memory of another processor. This quantity controls the size of MPI buffers but cannot be precisely determined a priori since it depends on the specific realization of the displacement field used in the simulation. We found for this specific test case that the memory overhead for the communication has a size of roughly 75Mb per processor.

The computational cost of our method is driven by the synthesis of the unlensed map which is the highest resolution object to be computed and has a number of pixels  $\kappa^2$  higher than the displacement field. As can be also seen from the right panel of 12, the runtime connected to the inverse spherical harmonic transform of the unlensed sky, despite perfectly load balanced, tends to flatten due to required internal communication steps and precomputations to initialize the recurrence to compute spherical harmonics. Those are per se subdominant steps but are expected to play a role in presence of a very fine parallelization (Szydlarski et al. 2011). The overall remapping procedure of pixel values require a number of operation of the order of  $O(N_{pix}/n)$  and is subdominant since it operates on a lower resolution map and perfectly scalable as it does not require any communication. The step in-

volving the reconstruction of of the lensed map after reshuffling of pixels (denoted as communication part in 12) is subdominant but the walltime required by this step is expected to slightly grow as it potentially involve the collective communication of small chunks of data between processors and is expected to approach the latency limit for message sending.

## 5. Conclusions

We have investigated and clarified details of modeling and simulations of the gravitational lensing effect on CMB. We particularly have aimed at elucidating the role and impact of band widths of considered signals on a precision of the pixel-domain approaches (e.g., Lewis 2005) to simulating lensing effect on polarized anisotropies, paying a special attention to a recovery of the  $B$  component. Such band widths play a crucial role in ensuring the sufficient accuracy of the produced lensed maps and need to be carefully taken into account if numerical effects such as power aliasing are to be kept under control. In this work we have developed a semi-analytic approach based on the formalism of Hu (2000) to study such effects and with its help investigated the necessary requirements for the signal's bandwidths, ensuring that sufficiently precise lensed signals can be recovered from such codes. In particular, we have found out that the simple convolution picture, where the convolution kernel, due to the grav-



itational potential, has a limited width of at most few hundred in  $\ell$  space, though working very well for the total intensity and  $E$  polarization up to  $\tilde{\ell}^{T/P} \lesssim 2000$ , is sufficient neither for much smaller angular scales in these two cases nor for the  $B$ -mode signal. Instead the proposed semi-analytic formalism should be used to guide a selection of the simulation parameters in order to ensure the final precision of the result but also to optimize the computational time.

We have demonstrated this via extensive numerical computations, for which we have developed a simple, massively parallel lensing simulation code, `lenS2HAT`. The code uses an extremely simple but robust approach to the interpolation involving sky overpixelization and a simple NGP assignment scheme, which, as we have shown, lead to easily understandable and controllable numerical effects. The latter effects are minimized as the code, thanks to its efficient parallelization, permits analyses of extremely large sky maps involving very dense sky grids/pixelization and therefore resolving the simulated sky power all the way to its actual bandwidths, which are carefully kept track of in the code.

The developed code, `lenS2HAT`, is suitable for massively parallel computational platforms, with either shared or distributed memory, and displays nearly perfect scalability in terms of the runtime and memory per processor up to the maximal number of CPUs it can employ. The latter is determined by the lowest value of the bandlimit parameters for either CMB or displacement field which is to be used in the runs,  $n_{proc}^{max} = \min(\ell_{max}^E, \ell_{max}^\Phi)/2$ . As is, it therefore permits extensive simulations involving hundreds of simulated maps in a reasonable time. The major bottleneck of the code performance is due to the need of calculating a single inverse spherical harmonic transform required to obtain the overpixelized map of the unlensed signal. This can be certainly alleviated further by using better algorithms and/or numerical implementations, e.g., capitalizing on hardware accelerators such as GPGPU (Hupca et al. 2012; Szydlarski et al. 2011; Fabbian et al. 2012). We leave this kind of code optimizations for future work. The code and its forthcoming version will be publicly available.

**Acknowledgements.** This work has been supported in part by French National Research Agency (ANR) through COSINUS program (project MIDAS no. ANR-09-COSI-009). GF acknowledges the support of Università degli Studi di Milano Bicocca and CARIPLO Foundation through the EXTRA program for the preliminary phase of this work. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We acknowledge use of the publicly available software tools and libraries HEALPIX, SPRNG, IPM, S<sup>2</sup>HAT, CAMB, and LensPix.

## References

- Basak, S., Prunet, S., & Benabed, K. 2009, *A&A*, 508, 53  
 Blanchard, A. & Schneider, J. 1987, *A&A*, 184, 1  
 Carbone, C., Springel, V., Baccigalupi, C., Bartelmann, M., & Matarrese, S. 2008, *MNRAS*, 388, 1618  
 Challinor, A. & Lewis, A. 2005, *Phys. Rev. D*, 71, 103010  
 Das, S. & Bode, P. 2008, *ApJ*, 682, 1  
 Das, S. & Linder, E. V. 2012, *Phys. Rev. D*, 86, 063520  
 Das, S., Sherwin, B. D., Aguirre, P., et al. 2011, *Physical Review Letters*, 107, 021301  
 de Putter, R., Zahn, O., & Linder, E. V. 2009, *Phys. Rev. D*, 79, 065033  
 Driscoll, J. R. & Healy, D. M. 1994, *Advances in Applied Mathematics*, 15, 202  
 Errard, J. & Stompor, R. 2012, *Phys. Rev. D*, 85, 083006  
 Fabbian, G., Szydlarski, M., Stompor, R., Grigori, L., & Falcou, J. 2012, in *Astronomical Society of the Pacific Conference Series*, Vol. 461, *Astronomical Data Analysis Software and Systems XXI*, ed. P. Ballester, D. Egret, & N. P. F. Lorente, 61  
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759  
 Hall, A. C. & Challinor, A. 2012, *MNRAS*, 425, 1170

- Hirata, C. M., Ho, S., Padmanabhan, N., Seljak, U., & Bahcall, N. A. 2008, *Phys. Rev. D*, 78, 043520  
 Hirata, C. M., Padmanabhan, N., Seljak, U., Schlegel, D., & Brinkmann, J. 2004, *Phys. Rev. D*, 70, 103501  
 Hu, W. 2000, *Phys. Rev. D*, 62, 043007  
 Hu, W. & Okamoto, T. 2002, *ApJ*, 574, 566  
 Hupca, I., Falcou, J., Grigori, L., & Stompor, R. 2012, in *Euro-Par 2011: Parallel Processing Workshops*, Springer, 355–366  
 Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, *Phys. Rev. D*, 55, 7368  
 Kermish, Z. D., Ade, P., Anthony, A., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8452, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*  
 Kesden, M., Cooray, A., & Kamionkowski, M. 2002, *Phys. Rev. Lett.*, 89, 011304  
 Lavaux, G. & Wandelt, B. D. 2010, *ApJS*, 191, 32  
 Lewis, A. 2005, *Phys. Rev. D*, 71, 083008  
 Lewis, A. & Challinor, A. 2006, *Phys. Rep.*, 429, 1  
 McMahon, J. J., Aird, K. A., Benson, B. A., et al. 2009, in *American Institute of Physics Conference Series*, Vol. 1185, *American Institute of Physics Conference Series*, ed. B. Young, B. Cabrera, & A. Miller, 511–514  
 Merkel, P. M. & Schäfer, B. M. 2011, *MNRAS*, 411, 1067  
 Niemack, M. D., Ade, P. A. R., Aguirre, J., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7741, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*  
 Okamoto, T. & Hu, W. 2003, *Phys. Rev. D*, 67, 083002  
 Oxley, P., Ade, P. A., Baccigalupi, C., et al. 2004, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 5543, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ed. M. Strojnik, 320–331  
 Planck Collaboration. 2013, *ArXiv e-prints*  
 Seljak, U. & Hirata, C. M. 2004, *Phys. Rev. D*, 69, 043005  
 Smith, K. M., Hanson, D., LoVerde, M., Hirata, C. M., & Zahn, O. 2012, *J. Cosmology Astropart. Phys.*, 6, 14  
 Smith, K. M., Zahn, O., & Doré, O. 2007, *Phys. Rev. D*, 76, 043510  
 Szydlarski, M., Esterie, P., Falcou, J., Grigori, L., & Stompor, R. 2011, *ArXiv e-prints*  
 Tomaru, T., Hazumi, M., Lee, A. T., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8452, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*  
 van Engelen, A., Keisler, R., Zahn, O., et al. 2012, *ApJ*, 756, 142  
 Zaldarriaga, M. & Seljak, U. 1997, *Phys. Rev. D*, 55, 1830  
 Zaldarriaga, M. & Seljak, U. 1998, *Phys. Rev. D*, 58, 023003

## Appendix A: Accuracy formula precision

We show below the computation demonstrating the precision of the accuracy definition. It has been show in literature (Challinor & Lewis 2005) that the harmonic approximation of Hu (2000) can reproduce lensed CMB spectra accurate to the percent level when compared to more sophisticated and accurate Boltzmann codes. Assuming we had access to the true 2d kernels, our accuracy formula relates to the true accuracy formula in the following way.

$$A_{\ell}^{X,true}(\ell^{\Phi}, \ell^X) = \frac{\alpha_{\ell^X, \ell^{\Phi}}^{\ell} + \zeta_{\ell^X, \ell^{\Phi}}}{\alpha_{\ell_{max}^X, \ell_{max}^{\Phi}}^X + \epsilon} \quad (A.1)$$

where we have recast the numerator and denominator of Eq. 20 in the two terms  $\alpha_{\ell^X, \ell^{\Phi}}^{\ell}$  and  $\alpha_{\ell_{max}^X, \ell_{max}^{\Phi}}^{\ell}$  respectively and we added to both an effective error term depending in general on the cutoff of the sum of both numerator and denominator of eq. 20. Being the fractional accuracy of harmonic expansion of the order of percent and assuming an absolute cutoff for CMB and lensing potential,  $\ell_{max}$ , sufficiently high, we can Taylor expand the last expression to first order as

$$A_{\ell}^{X,true}(\ell^{\Phi}, \ell^X) \approx \frac{(\alpha_{\ell^X, \ell^{\Phi}}^{\ell} + \zeta_{\ell^X, \ell^{\Phi}})}{\alpha_{\ell_{max}^X, \ell_{max}^{\Phi}}^{\ell}} (1 - \beta) + O(\beta^2) \quad (A.2)$$

$$= A_{\ell}^X(\ell^{\Phi}, \ell^X) - \beta A_{\ell}^X(\ell^{\Phi}, \ell^X) + \frac{\zeta_{\ell^X, \ell^{\Phi}}}{\alpha_{\ell_{max}^X, \ell_{max}^{\Phi}}}$$

where we called  $\beta$  the fractional error between harmonic method and the true value. Calling  $\delta$  the fractional difference between the effective error in the nominator and the  $\alpha_{\ell_{max}^X, \ell_{max}^{\Phi}}$  term, the precision an the accuracy function then becomes

$$\left| \frac{\Delta A_{\ell}^X(\ell^{\Phi}, \ell^X)}{A_{\ell}^X(\ell^{\Phi}, \ell^X)} \right| \approx \left| \frac{1}{A_{\ell}^X(\ell^{\Phi}, \ell^X)} (-\beta A_{\ell}^X(\ell^{\Phi}, \ell^X) + \delta_{\ell^X, \ell^{\Phi}}) \right| \quad (\text{A.3})$$

$$\approx (-\beta + \delta_{\ell^X, \ell^{\Phi}}) + O(\beta^2)$$

where we have considered our accuracy function as a first order expansion around one. This assumption is valid in the regimes where the accuracy function assumes values close to one, which is indeed the case for values of accuracy close or smaller than one percent and thus a first order series expansion is well within the linear order we considered here.

The overall precision of the accuracy function is then driven by the difference between the two terms which by construction tends to compensate since  $\zeta_{\ell^X, \ell^{\Phi}}$  tends to  $\epsilon$  as  $\ell^X, \ell^{\Phi}$  approach the  $\ell_{max}$  value. The formula then becomes more and more accurate as we approach the cutoff limit.

## Appendix B: lenS<sup>2</sup>HAT code outline

The code outline follows the general simulation guidelines outlines in Sect. 2.2.1 but we explicit here a few details of potential interest of the code's structure.

1. When generating a Gaussian random realization of harmonic coefficient for unlensed CMB and displacement field, both the correlation between temperature or E modes and displacement field due to Sachs-Wolfe effect can be taken into account if requested. However since both are negligible for most of the multipoles, we neglected it in the runs performed for this paper. We do not expect this correlation to affect the results of our analysis especially in the high  $\ell$  tail of the spectrum given that the correlation is confined to large angular scales.
2. The non linear corrections in LSS evolution are naturally taken into account in the code if those are included in an effective input lensing potential power spectrum. Even though non linear evolution of perturbation induces non-gaussianities in the matter power spectrum, the contribution of higher order moment to the lensing potential have been proven to be on the sub percent level (Merkel & Schäfer 2011) and thus the assumption of a purely gaussian lensing potential is well applicable.
3. Since harmonic coefficients are distributed between processors and generated directly in a distributed way, we use the Scalable Parallel Random Number Generator library<sup>9</sup> (SPRNG) to avoid correlation between random number streaming on each processors.
4. The computation of displaced coordinates and remapping of pixel location is performed by two separate routines, one optimized for pixelizations having equidistant rings

(e.g. ECP) while the other suitable for any pixelization conforming to the requirements of symmetries set by S<sup>2</sup>HAT. Since it can happen that such a pixel doesn't belong to the set of pixel stored in the local memory of a single processor, we identify it through its index in a global indexing scheme and store this global index in a bi-dimensional array having dimension equal to the number of processors used times number of pixel stored on that particular processor, where the line index represent thus the rank of processor.

5. A collective `MPI_All2allv` communication step is performed to redistribute information of pixels which are needed by processors but are not stored in their local memory to build the final lensed map. This pattern ensures an even distribution of memory between all cores and a very good scalability up to several thousands MPI processes. On the numerical level the communication time is subdominant but can in principle be further optimized with non blocking MPI local communication calls or exploiting an hybrid MPI/OpenMP approach.
6. The code can perform simulation in every pixelization scheme although we found that ECP has to be preferred for the computation and the results can later be resampled of the kind of grid of interest.
7. The code supports simultaneous multi map analysis on the spherical harmonics step of the algorithm. If memory specification of the architecture in use allow to fit several maps in a memory of a single processor, a cutoff of a factor of number of maps in the runtime is expected when processing up to few maps at the same time. The latter option make the code particularly appealing for any analysis or simulation step requiring the use of Montecarlo methods.

<sup>9</sup> <http://sprng.cs.fsu.edu>